Doomsday Equilibria for Omega-Regular Games

Krishnendu Chatterjee[a,*], Laurent Doyen[b], Emmanuel Filiot[c,**], Jean-François Raskin[c,]

[a] *IST Austria*
[b] *LSV, ENS Cachan & CNRS*
[c] *CS, Université libre de Bruxelles (U.L.B.)*

Abstract

Two-player games on graphs provide the theoretical framework for many important problems such as reactive synthesis. While the traditional study of two-player zero-sum games has been extended to multi-player games with several notions of equilibria, they are decidable only for perfect-information games, whereas several applications require imperfect-information games.

In this paper we propose a new notion of equilibria, called doomsday equilibria, which is a strategy profile such that all players satisfy their own objective, and if any coalition of players deviates and violates even one of the players objective, then the objective of every player is violated.

We present algorithms and complexity results for deciding the existence of doomsday equilibria for various classes of ω-regular objectives, both for imperfect-information games, and for perfect-information games. We provide optimal complexity bounds for imperfect-information games, and in most cases for perfect-information games.

1. Introduction

Two-player games on finite-state graphs with ω-regular objectives provide the framework to study many important problems in computer science [37, 34, 15]. One key application area is synthesis of reactive systems [7, 35, 33]. Traditionally, the reactive synthesis problem is reduced to two-player zero-sum games, where vertices of the graph represent states of the system, edges represent transitions, one player represents a component of the system to synthesize, and the other player represents the purely adversarial coalition of all the other components. Since the coalition is adversarial, the game is zero-sum, i.e., the objectives of the two players are complementary. Two-player zero-sum games have been studied in great depth in the literature [27, 15, 22].

Instead of considering all the other components as purely adversarial, a more realistic model is to consider them as individual players each with their own objective, as in protocol synthesis where the rational behavior of the agents is to first satisfy their own objective in the protocol before trying to be adversarial to the other agents. Hence, inspired by recent applications in protocol synthesis, the model of multi-player games on graphs has become an active area of research in graph games and reactive synthesis [2, 21, 38]. In a multi-player setting, the games are not necessarily zero-sum (i.e., objectives are not necessarily conflicting) and the classical notion of rational behavior is formalized as Nash equilibria [30]. Nash equilibria perfectly capture the notion of rational behavior in the absence of external criteria, i.e., the players are concerned only about their own payoff (internal criteria), and they are indifferent to the payoff of the other players. In the setting

[☆] A short version of this paper has been published by the same authors in the proceedings of VMCAI'14. This version contains the full proofs of the results presented in the proceedings version as well as new additional results that cover games with LTL objectives.

[*] Supported by Austrian Science Fund (FWF) Grant No P23499-N23, FWF NFN Grant No S11407-N23 (RiSE), ERC Start grant (279307: Graph Games), and Microsoft faculty fellows award.

[**] Supported by the Belgian National Fund for Scientific Research.

Supported by ERC Start grant (279499: inVEST).

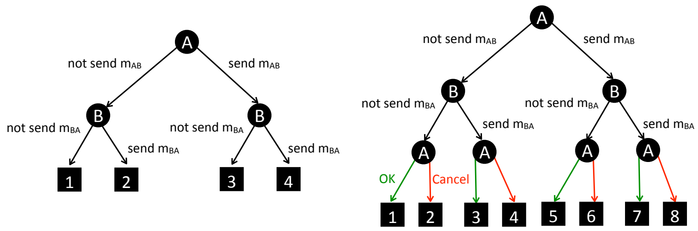

Figure 1: A simple example in the domain of Fair Exchange Protocols

of synthesis, the more appropriate notion is the adversarial external criteria, where the players are as harmful as possible to the other players without sabotaging with their own objectives. This has inspired the study of refinements of Nash equilibria, such as secure equilibria (that captures the adversarial external criteria), and rational synthesis, and led to several new logics where the non-zero-sum equilibria can be expressed. The complexity of Nash equilibria [38], secure equilibria [11], rational synthesis [21], and of the new logics has been studied recently [12, 14, 29, 39, 28].

Along with the theoretical study of refinements of equilibria, applications have also been developed in the synthesis of protocols. In particular, the notion of secure equilibria has been useful in the synthesis of mutual-exclusion protocol [11], and of fair-exchange protocols [25, 9] (a key protocol in the area of security for exchange of digital signatures). One major drawback that all the notions of equilibria suffer is that the basic decision questions related to them are decidable only in the setting of perfect-information games (in a perfect-information games the players perfectly know the state and history of the game, whereas in imperfect-information games each player has only a partial view of the state space of the game), and in the setting of multi-player imperfect-information games they are undecidable [33]. However, the model of imperfect-information games is very natural because every component of a system has private variables not accessible to other components, and recent works have demonstrated that imperfect-information games are required in synthesis of fair-exchange protocols [24]. In this paper, we provide the first decidable framework that can model them.

We propose a new notion of equilibria that we call *doomsday-threatening* equilibria (for short, doomsday equilibria). A doomsday equilibrium is a strategy profile such that all players satisfy their own objective, and if any coalition of players deviates and violates even one of the players objective, then doomsday follows (every player objective is violated). Note that in contrast to other notions of equilibria, doomsday equilibria consider deviation by an arbitrary set of players, rather than individual players. Moreover, in case of two-player non-zero-sum games they coincide with the secure equilibria [11] where objectives of both players are satisfied.

Example 1. Let us consider the two trees of Fig. 1. They model the possible behaviors of two entities Alice and Bob that have the objective of exchanging messages: $\mathsf{m_{AB}}$ from Alice to Bob, and $\mathsf{m_{BA}}$ from Bob to Alice. Assume for the sake of illustration that $\mathsf{m_{AB}}$ models the transfer of property of a house from Alice to Bob, while $\mathsf{m_{BA}}$ models the payment of the price of the house from Bob to Alice.

Having that interpretation in mind, let us consider the left tree. On the one hand, Alice has as primary objective (internal criterion) to reach either state 2 or state 4, states in which she has obtained the money, and she has a slight preference for 2 as in that case she received the money while not transferring the property of her house to Bob, this corresponds to her adversarial external criterion. On the other hand, Bob would like to reach either state 3 or 4 (with again a slight preference for 3). Also, it should be clear that Alice would hate to reach 3 because she would have transferred the property of her house to Bob but without being paid. Similarly, Bob would hate to reach 2. To summarize, Alice has the following preference order on the final states of the protocol: $2 > 4 > 1 > 3$, while for Bob the order is $3 > 4 > 1 > 2$. Is there a *doomsday threatening equilibrium* in this game ? For such an equilibrium to exist, we must find a pair of strategies that please the two players for their primary objective (internal criterion): reach $\{2, 4\}$ for Alice and reach $\{3, 4\}$ for Bob. Clearly, this is only possible if at the root Alice plays "send $\mathsf{m_{AB}}$", as otherwise we would

objectives	safety	reachability	Büchi	co-Büchi	parity	LTL
perfect info.	PSpace-C	PTime-C	PTime-C	PTime-C	PSpace NP-Hard CoNP-Hard	2ExpTime-C
imperfect info.	ExpTime-C	ExpTime-C	ExpTime-C	ExpTime-C	ExpTime-C	2ExpTime-C

Table 1: Summary of the results.

not reach $\{1, 2\}$ violating the primary objective of Bob. But playing that action is not safe for Alice as Bob would then choose "not send m_{BA}" because he slightly prefers 3 to 4. It can be shown that the only rational way of playing (taking into account both internal and external criteria) is for Alice to play "not send m_{AB}" and for Bob to play "not send m_{BA}". This way of playing is in fact the only secure equilibrium of the game but this is not what we hope from such a protocol.

The difficulty in this exchange of messages comes from the fact that Alice is starting the protocol by sending her part and this exposes her. To obtain a better behaving protocol, one solution is to add an additional stage after the exchanges of the two messages as depicted in the right tree of Fig. 1. In this new protocol, Alice has the possibility to cancel the exchange of messages (in practice this would be implemented by the intervention of a TTP[1]). For that new game, the preference orderings of the players are as follows: for Alice it is $3 > 7 > 1 = 2 = 4 = 6 = 8 > 5$, and for Bob it is $5 > 7 > 1 = 2 = 4 = 6 = 8 > 3$. Let us now show that there is a doomsday equilibrium in this new game. In the first round, Alice should play "send m_{AB}" as otherwise the internal objective of Bob would be violated, then Bob should play "send m_{BA}", and finally Alice should play "OK" to validate the exchange of messages. Clearly, this profile of strategies satisfies the first property of a doomsday equilibrium: both players have reached their primary objective. Second, let us show that no player has an incentive to deviate from that profile of strategies. First, if Alice deviates then Bob would play "not send m_{BA}", and we obtain a doomsday situation as both players have their primary objectives violated. Second, if Bob deviates by playing "not send m_{BA}", then Alice would cancel the protocol exchange which again produces a doomsday situation. So, no player has an incentive to deviate from the equilibrium and the outcome of the protocol is the desired one: the two messages have been fairly exchanged. So, we see that the threat of a doomsday brought by the action "Cancel" has a beneficial influence on the behavior of the two players. □

It should now be clear that multi-player games with doomsday equilibria provide a suitable framework to model various problems in protocol synthesis. In addition to the definition of doomsday equilibria, our main contributions are to present algorithms and complexity bounds for deciding the existence of such equilibria for various classes of ω-regular objectives both in the perfect-information and in the imperfect-information cases. In all cases but one, we establish the exact complexity. Our technical contributions are summarized in Table 1. More specifically:

1. *(Perfect-information games).* We show that deciding the existence of doomsday equilibria in multi-player perfect-information games is (i) PTime-C for reachability, Büchi, and coBüchi objectives; (ii) PSpace-C for safety objectives; (iii) in PSpace and both NP-Hard and coNP-Hard for parity objectives; (iv) 2ExpTime-C for LTL objectives.
2. *(Imperfect-information games).* We show that deciding the existence of doomsday equilibria in multi-player imperfect-information games is (i) ExpTime-C for reachability, safety, Büchi, coBüchi, and parity objectives, and (ii) it remains 2ExpTime-C for LTL objectives.

The research area of multi-player games has been quite active recently, but so far notions of equilibria that lead to decidability in the imperfect-information setting and have applications in synthesis have been largely unexplored. Our work is a step in that direction.

[1]TTP stands for *Trusted Third Party*.

2. Doomsday Equilibria for Perfect Information Games

In this section, we define game arena with perfect information, ω-regular objectives, and doomsday equilibria. We also recall classical definitions from automata theory and temporal logic.

Automata over infinite words. Let A be a finite alphabet. An *infinite word* over A is an infinite sequence $w = a_0 a_1 \dots a_n \dots$ where $a_i \in A$ for all $i \geq 0$. We denote the set of infinite words over A by A^ω and a subset of A^ω is called a *language* of infinite words. To define languages, we use finite automata. A finite automaton is a tuple $\mathcal{A} = (Q, q_{\mathsf{init}}, A, \delta)$ where Q is a finite set of states, q_{init} is the initial state, A is a finite alphabet, and $\delta : Q \times A \to 2^Q$ is the transition relation. $\mathcal{A}$ is *deterministic* if for all $q \in Q$ and $a \in A$, $|\delta(q, a)| \leq 1$, and it is *total* if for all $q \in Q$ and $a \in A$, $|\delta(q, a)| \geq 1$.

A *run* of $\mathcal{A}$ over a word $w = a_0 a_1 \dots a_n \dots$ is an infinite sequence of states $r = q_0 q_1 \dots q_n$ such that (a) $q_0 = q_{\mathsf{init}}$, (b) $(q_i, a_i, q_{i+1}) \in \delta$, for all $i \geq 0$. We denote by $\mathsf{inf}(r)$ the set of states q that appear infinitely often in r.

To define the language of $\mathcal{A}$ from the notion of run, we need to choose between the *existential* (nondeterministic) or the *universal* interpretation of the transition relation, and to fix an *acceptance condition.* For the existential interpretation, w is accepted by $\mathcal{A}$ if there exists at least *one* run on w that is accepting, and for the universal interpretation, w is accepted by $\mathcal{A}$ if *all* runs on w are accepting. We use several notions of acceptance:

- a *Büchi* condition is defined by a subset of accepting states $F \subseteq Q$, and a run $r = q_0 q_1 \dots q_n \dots$ is accepting if there are infinitely many positions $i \geq 0$ such that $q_i \in F$, i.e. $\mathsf{Inf}(r) \cap F \neq \emptyset$;

- a *coBüchi* condition is defined by a subset of accepting states $F \subseteq Q$, and a run $r = q_0 q_1 \dots q_n \dots$ is accepting if there are only finitely many positions $i \geq 0$ such that $q_i \in F$, i.e. $\mathsf{Inf}(r) \cap F = \emptyset$;

- a *Rabin* condition is defined by a set of pairs of subsets of Q, i.e. $\Omega = \{(F_1, G_1), (F_2, G_2), \dots, (F_m, G_m)\}$, and a run r is accepting if for all pair $(F, G) \in \Omega$, if $\mathsf{Inf}(r) \cap F = \emptyset$ then $\mathsf{Inf}(r) \cap G = \emptyset$;

- a *Streett* condition is defined by a set of pairs of subsets of Q, i.e. $\Omega = \{(F_1, G_1), (F_2, G_2), \dots, (F_m, G_m)\}$, and a run r is accepting if for all pair $(F, G) \in \Omega$, if $\mathsf{Inf}(r) \cap F \neq \emptyset$ then $\mathsf{Inf}(r) \cap G \neq \emptyset$. So this is the dual of the Rabin condition;

- a *Parity* condition is defined by a function $p : Q \to \{0, 1, \dots, d\}$ and a run r is accepting if the smallest priority visited infinitely along r is even: $\mathsf{parity}(p) = \{r \in Q^\omega \,|\, \min\{p(q) \mid q \in \mathsf{inf}(r)\} \text{ is even}\}$.

In this paper, we use nondeterministic Büchi automata, NBW for short, universal co-Büchi automata, UcoBW for short, deterministic parity automata, DPW for short, and determinisitc Streett automata, DSW for short. Note that for deterministic and total automata, the universal and existential interpretation coincide.

Game Arena. An n-player game arena G with perfect information is defined as a tuple $(S, \mathcal{P}, s_{\mathsf{init}}, \Sigma, \Delta)$ such that S is a nonempty finite set of *states*, $\mathcal{P} = \{S_1, S_2, \dots, S_n\}$ is a partition of S into n classes of states, one for each player respectively, $s_{\mathsf{init}} \in S$ is the initial state, Σ is a finite set of actions, and $\Delta : S \times \Sigma \to S$ is the transition function.

Plays in n-player game arena G are constructed as follows. They start in the initial state s_{init}, and then an ω number of rounds are played: the player that owns the current state s chooses a letter $\sigma \in \Sigma$ and the game evolves to the position $s' = \Delta(s, \sigma)$, then a new round starts from s'. So formally, a *play* in G is an infinite sequence $s_0 s_1 \dots s_n \dots$ such that (a) $s_0 = s_{\mathsf{init}}$ and (b) for all $i \geq 0$, there exists $\sigma \in \Sigma$ such that $s_{i+1} = \Delta(s_i, \sigma)$. The set of plays in G is denoted by $\mathsf{Plays}(G)$, and the set of finite prefixes of plays by $\mathsf{PrefPlays}(G)$. We denote by $\rho, \rho_1, \rho_i, \dots$ plays in G, by $\rho(0..j)$ the prefix of the play ρ up to position j and by $\rho(j)$ the position j in the play ρ. We also use π, π_1, π_2, $\dots$ to denote prefixes of plays. Let $i \in \{1, 2, \dots, n\}$, a prefix π belongs to Player i if $\mathsf{last}(\pi)$, the last state of π, belongs to Player i, i.e. $\mathsf{last}(\pi) \in S_i$. We denote by $\mathsf{PrefPlays}_i(G)$ the set of prefixes of plays in G that belongs to Player i.

To interpret the truth value of LTL formulas over plays, we need the notion of a *labeled* game arena. An n-player labeled game arena G with perfect information is a tuple $(S, \mathcal{P}, s_{\mathsf{init}}, \Sigma, \Delta, \mathbb{P}, \mathcal{L})$ where $(S, \mathcal{P}, s_{\mathsf{init}}, \Sigma, \Delta)$

is a game arena, $\mathbb{P}$ is a set of atomic propositions, and $\mathcal{L} : S \to 2^{\mathbb{P}}$ is a labelling function that assigns to each state $s \in S$, the subset of atomic propositions $\mathcal{L}(s)$ that hold in s (the complement of this set is the set of atomic propositions that do not hold in s). Given a play $\rho = s_0 s_1 \dots s_n \dots$, we write $\mathcal{L}(\rho)$ for the sequence of corresponding labels, i.e. $\mathcal{L}(\rho) = \mathcal{L}(s_0)\mathcal{L}(s_1) \dots \mathcal{L}(s_n) \dots$, and call $\mathcal{L}(\rho)$ a *labelled play* of G. The syntax of LTL formulas over $\mathbb{P}$ is as follows:

$$\phi ::= p \mid \phi_1 \vee \phi_2 \mid \neg\phi \mid \phi_1 \mathcal{U} \phi_2$$

with the classical semantics, a formula ϕ evaluates to true on a labelled play $\mathcal{L}(\rho)$, noted $\mathcal{L}(\rho) \models \phi$, according to the following rules:

- $\mathcal{L}(\rho) \models \phi$ iff $\mathcal{L}(\rho), 0 \models \phi$, and for all $i \geq 0$,
- $\mathcal{L}(\rho), i \models p$ iff $p \in \mathcal{L}(\rho(i))$
- $\mathcal{L}(\rho), i \models \phi_1 \vee \phi_2$ iff $\mathcal{L}(\rho), i \models \phi_1$ or $\mathcal{L}(\rho), i \models \phi_2$
- $\mathcal{L}(\rho), i \models \phi_1 \mathcal{U} \phi_2$ iff there exists $j \geq i$, $\mathcal{L}(\rho), j \models \phi_2$, and for all k, $i < k < j$, $\mathcal{L}(\rho), k \models \phi_1$

As usual, we use abbreviations: $\top \equiv \phi \vee \neg\phi$, $\Diamond\phi \equiv \top\mathcal{U}\phi$, $\Box\phi \equiv \neg\Diamond\neg\phi$, and $\bigcirc\phi \equiv \bot\mathcal{U}\phi$.

Strategies and strategy profiles. A *strategy* for Player i, for $i \in \{1, 2, \dots, n\}$, is a mapping $\lambda_i : \mathsf{PrefPlays}_i(G) \to \Sigma$ from prefixes of plays to actions. A *strategy profile* $\Lambda = (\lambda_1, \lambda_2, \dots, \lambda_n)$ is a tuple of strategies such that λ_i is a strategy of Player i. The strategy of Player i in Λ is denoted by Λ_i, and the tuple of the remaining strategies $(\lambda_1, \dots, \lambda_{i-1}, \lambda_{i+1}, \dots, \lambda_n)$ by Λ_{-i}. For a strategy λ_i of Player i, we define its *outcome* as the set of plays that are consistent with λ_i: formally, $\mathsf{outcome}_i(\lambda_i)$ is the set of $\rho \in \mathsf{Plays}(G)$ such that for all $j \geq 0$, if $\rho(0..j) \in \mathsf{PrefPlays}_i(G)$, then $\rho(j+1) = \Delta(\rho(j), \lambda_i(\rho(0..j)))$. Similarly, we define the *outcome of a strategy profile* $\Lambda = (\lambda_1, \lambda_2, \dots, \lambda_n)$, as the unique play $\rho \in \mathsf{Plays}(G)$ such that for all positions j, for all $i \in \{1, 2, \dots, n\}$, if $\rho(j) \in \mathsf{PrefPlays}_i(G)$ then $\rho(j+1) = \Delta(\rho(j), \lambda_i(\rho(0..j)))$. Finally, given a state $s \in S$ of the game, we denote by G_s the game G whose initial state is replaced by s.

Winning objectives. A *winning objective* (or an *objective* for short) φ_i for Player $i \in \{1, 2, \dots, n\}$ is a set of infinite sequences of states, i.e. $\varphi_i \subseteq S^\omega$. A strategy λ_i is *winning* for Player i (against all other players) w.r.t. an objective φ_i if $\mathsf{outcome}_i(\lambda_i) \subseteq \varphi_i$.

Given an infinite sequence of states $\rho \in S^\omega$, we denote by $\mathsf{visit}(\rho)$ the set of states that appear at least once along ρ, i.e. $\mathsf{visit}(\rho) = \{s \in S | \exists i \geq 0 \cdot \rho(i) = s\}$, and $\mathsf{inf}(\rho)$ the set of states that appear infinitely often along ρ, i.e. $\mathsf{inf}(\rho) = \{s \in S | \forall i \geq 0 \cdot \exists j \geq i \cdot \rho(j) = s\}$. We consider the following types of winning objectives:

- a *safety objective* is defined by a subset of states $T \subseteq S$ that has to be never left: $\mathsf{safe}(T) = \{\rho \in S^\omega \mid \mathsf{visit}(\rho) \subseteq T\} = T^\omega$;
- a *reachability objective* is defined by a subset of states $T \subseteq S$ that has to be reached: $\mathsf{reach}(T) = \{\rho \in S^\omega \mid \mathsf{visit}(\rho) \cap T \neq \emptyset\}$;
- a *Büchi objective* is defined by a subset of states $T \subseteq S$ that has to be visited infinitely often: $\mathsf{Büchi}(T) = \{\rho \in S^\omega \mid \mathsf{inf}(\rho) \cap T \neq \emptyset\}$;
- a *co-Büchi objective* is defined by a subset of states $T \subseteq S$ that has to be reached eventually and never be left: $\mathsf{coBüchi}(T) = \{\rho \in S^\omega \mid \mathsf{inf}(\rho) \subseteq T\}$;
- let $d \in \mathbb{N}$, a *parity objective with d priorities* is defined by a priority function $p : S \to \{0, 1, \dots, d\}$ as the set of plays such that the smallest priority visited infinitely often is even: $\mathsf{parity}(p) = \{\rho \in S^\omega | \min\{p(s) \mid s \in \mathsf{inf}(\rho)\} \text{ is even}\}$;
- an *LTL objective* is defined for a $\mathbb{P}$-labelled game arena by an LTL formula ϕ over the set of atomic propositions $\mathbb{P}$, $[\![\phi]\!] = \{\rho \in S^\omega | \mathcal{L}(\rho) \models \phi\}$.

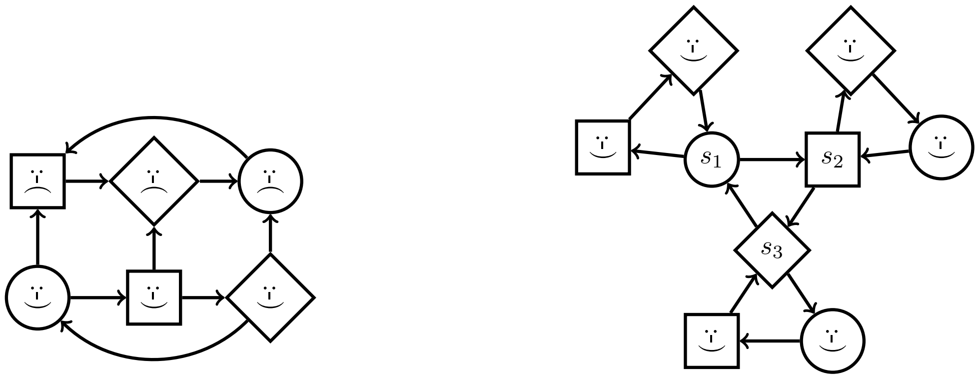

(a) Doomsday (Safety) (b) Büchi objectives

Figure 2: Examples of doomsday equilibria for Safety and Büchi objectives. We use circles, boxes and diamonds respectively, to depict states of Player 1, Player 2, and Player 3 respectively.

Büchi, co-Büchi and parity objectives φ are called *tail objectives* because they satisfy the following closure property: for all $\rho \in S^\omega$ and all $\pi \in S^*$, $\rho \in \varphi$ iff $\pi \cdot \rho \in \varphi$.

Finally, given an objective $\varphi \subseteq S^\omega$ and a subset $P \subseteq \{1, \ldots, n\}$, we write $\langle\langle P \rangle\rangle \varphi$ to denote the set of states s from which the players from P can cooperate to enforce φ when they start playing in s. Formally, $\langle\langle P \rangle\rangle \varphi$ is the set of states s such that there exists a set of strategies $\{\lambda_i \mid i \in P\}$ in G_s, one for each player in P, such that $\bigcap_{i \in P} \mathsf{outcome}_i(\lambda_i) \subseteq \varphi$.

Doomsday Equilibria. A strategy profile $\Lambda = (\lambda_1, \lambda_2, \ldots, \lambda_n)$ is a *doomsday-threatening equilibrium* (doomsday equilibrium or DE for short) if:

1. it is winning for all the players, i.e. $\mathsf{outcome}(\Lambda) \in \bigcap_i \varphi_i$;
2. each player is able to retaliate in case of deviation: for all $1 \leq i \leq n$, for all $\rho \in \mathsf{outcome}_i(\lambda_i)$, if $\rho \notin \varphi_i$, then $\rho \in \bigcap_{j=1}^{j=n} \overline{\varphi_j}$ (doomsday), where $\overline{\varphi_j}$ denotes the complement of φ_j in S^ω.

In other words, when all players stick to their strategies then they all win, and if any arbitrary coalition of players deviates and makes even just one other player lose then this player retaliates and ensures a doomsday, i.e. all players lose.

Relation with Secure Equilibria. In two-player games, the doomsday equilibria coincide with the notion of secure equilibrium [11] where both players satisfy their objectives. In secure equilibria, for all $i \in \{1, 2\}$, any deviation of Player i that does not decrease her payoff does not decrease the payoff of Player $3-i$ either. In other words, if a deviation of Player i decreases (strictly) the payoff of Player $3-i$, i.e. φ_{3-i} is not satisfied, then it also decreases her own payoff, i.e. φ_i is not satisfied. A two-player secure equilibrium where both players satisfy their objectives is therefore a doomsday equilibrium.

Example 2. Fig. 2 gives two examples of games with safety and Büchi objectives respectively. Actions are in bijection with edges so they are not represented.

(Safety) Consider the 3-player game arena with perfect information of Fig. 2(a) and safety objectives. Unsafe states for each player are given by the respective nodes of the upper part. Assume that the initial state is one of the safe states. This example models a situation where three countries are in peace until one of the countries, say country i, decides to attack country j. This attack will then necessarily be followed by a doomsday situation: country j has a strategy to punish all other countries. The doomsday equilibrium in this example is to play safe for all players.

(Büchi) Consider the 3-player game arena with perfect information of Fig. 2(b) with Büchi objectives for each player: Player i wants to visit infinitely often one of its "happy" states. The position of the initial state does not matter. To make things more concrete, let us use this game to model a protocol where 3 players want to share in each round a piece of information made of three parts: for all $i \in \{1, 2, 3\}$, Player i knows information $i \ mod \ 3 + 1$ and $i \ mod \ 3 + 2$. Player i can send or not these pieces of information to the other players. This is modeled by the fact that Player i can decide to visit the happy states of the other players, or move directly to $s_{(i \ mod \ 3)+1}$. The objective of each player is to have an infinite number of successful rounds where they get all the information.

There are several doomsday equilibria. As a first one, let us consider the situation where for all $i \in \{1, 2, 3\}$, if Player i is in state s_i, first it visits the happy states, and when the play comes back in s_i, it moves to $s_{(i \ mod \ 3)+1}$. This defines an infinite play that visits all the states infinitely often. Whenever some player deviates from this play, the other players retaliate by always choosing in the future to go to the next s state instead of taking their respective loops. Clearly, if all players follow their respective strategies all happy states are visited infinitely often. Now consider the strategy of Player i against two strategies of the other players that makes him lose. Clearly, the only way Player i loses is when the two other players eventually never take their states, but then all the players lose.

As a second one, consider the strategies where Player 2 and Player 3 always take their loops but Player 1 never takes his loop, and such that whenever the play deviates, Player 2 and 3 retialate by never taking their loops. For the same reasons as before this strategy profile is a doomsday equilibrium.

Note that the first equilibrium requires one bit of memory for each player, to remember if they visit their s state for the first or second times. In the second equilibrium, only Player 2 and 3 needs a bit of memory. An exhaustive analysis shows that there is no memoryless doosmday equilibrium in this example. ☐

3. Complexity of DE for Perfect Information Games

In this section, we prove the following results:

Theorem 3. *The problem of deciding the existence of a doomsday equilibrium in an n-player perfect information game arena and n objectives $(\varphi_i)_{1 \le i \le n}$ is:*

- PTime-C *if the objectives $(\varphi_i)_{1 \le i \le n}$ are either all Büchi, all co-Büchi or all reachability objectives, and hardness already holds for 2-player game arenas,*
- NP-hard, coNP-hard *and in* PSPace *if $(\varphi_i)_{1 \le i \le n}$ are parity objectives, and hardness already holds for 2-player game arenas,*
- PSPace-C *if $(\varphi_i)_{1 \le i \le n}$ are safety objectives, and* PTime-C *for game arenas with a fixed number of players,*
- 2ExpTime-C *if $(\varphi_i)_{1 \le i \le n}$ are LTL objectives, and hardness already holds for 2-player game arenas.*

The remainder of this section gives detailed proofs for those results. In the sequel, game arena with perfect information are just called game arena.

3.1. Tail objectives

We first present a generic algorithm that works for any tail objective and then analyze its complexity for the different cases. Then we establish the lower bounds. Let us consider the following algorithm:

- compute the retaliation region of each player: $R_i = \langle\langle i \rangle\rangle(\varphi_i \cup \bigcap_{j=1}^{j=n} \overline{\varphi_j})$;
- check for the existence of a play within $\bigcap_{i=1}^{i=n} R_i$ that satisfies all the objectives φ_i.

The correctness of this generic procedure is formalized in the following lemma:

Lemma 4. *Let $G = (S, \mathcal{P}, s_{\mathsf{init}}, \Sigma, \Delta)$ be an n-player game arena with n tail objectives $(\varphi_i)_{1 \le i \le n}$. Let $R_i = \langle\langle i \rangle\rangle(\varphi_i \cup \bigcap_{j=1}^{j=n} \overline{\varphi_j})$ be the retaliation region for Player i. There is a doomsday equilibrium in G iff there exists an infinite play that* (1) *belongs to $\bigcap_{i=1}^{i=n} \varphi_i$ and* (2) *belongs to the set of states $\bigcap_{i=1}^{i=n} R_i$.*

Proof. First, assume that there exists an infinite play ρ such that $\rho \in \bigcap_i(\varphi_i \cap R_i^\omega)$. From ρ, and the retaliating strategies that exist in all states of R_i for each player, we show the existence of a DE $\Lambda = (\lambda_1, \lambda_2, \dots, \lambda_n)$. Player i plays strategy λ_i as follows: he plays according to the choices made in ρ as long as all the other players do so, and as soon as the play deviates from ρ, Player i plays his retaliating strategy. Note that if the turn never comes back to Player i then we know that the play always stays within R_i and all outcomes satisfy $\varphi_i \cup \bigcap_{j=1}^{j=n} \overline{\varphi_j}$.

First, let us show that if Player j, for some $j \neq i$, deviates and the turn comes back to Player i in a state s then $s \in R_i$. Assume that Player j deviates when he is in some $s' \in S_j$. As before there was no deviation, by definition of ρ, s' belongs to R_i. But no matter what the adversaries are doing in a state that belongs to R_i, the next state must be a state that belongs to R_i (there is only the possibility to leave R_i when Player i plays). So, by induction on the length of the segment of play that separates s' and s, we can conclude that s belongs to R_i. From s, Player i plays a retaliating strategy and so all the outcomes from s are in $\varphi_i \cup \bigcap_{j=1}^{j=n} \overline{\varphi_j}$, and since tail objectives are closed under complement, the prefix up to s is not important and we get (from s_{init}) $\mathsf{outcome}_i(\lambda_i) \subseteq \varphi_i \cup \bigcap_{j=1}^{j=n} \overline{\varphi_j}$. Therefore the second property of the definition of doomsday equilibria is satisfied. Hence Λ is a DE.

Let us now consider the other direction. Assume that Λ is a DE. Then let us show that $\rho = \mathsf{outcome}(\Lambda)$ satisfies properties (1) and (2). By definition of DE, we know that ρ is winning for all the players, so (1) is satisfied. Again by definition of DE, $\mathsf{outcome}(\Lambda_i) \subseteq \varphi_i \cup \bigcap_{j=1}^{j=n} \overline{\varphi_j}$. Let s be a state of ρ and π the prefix of ρ up to s. For all outcomes ρ' of Λ_i in G_s, we have $\pi\rho' \in \varphi_i \cup \bigcap_{j=1}^{j=n} \overline{\varphi_j}$, and since all φ_i are tail objectives, and Boolean combinations of tail objectives are tail objectives, we get $\rho' \in \varphi_i \cup \bigcap_{j=1}^{j=n} \overline{\varphi_j}$. Hence $s \in R_i$. Since this property holds for all i, we get $s \in \bigcap_i R_i$, and (2) is satisfied. $\square$

Upper bounds. Accordingly, we obtain the following upper bounds:

Lemma 5. *The problem of deciding the existence of a doomsday equilibrium in an n-player game arena can be solved in* PTime *for Büchi and co-Büchi objectives, and in* PSpace *for parity objectives.*

Proof. By Lemma 4, one first needs to compute the retaliation regions R_i for all $i \in \{1, \dots, n\}$. Once the sets R_i have been computed, it is clear that the existence of a play winning for all players is decidable in PTime for all the three types of objective. We detail the cost of computing R_i for each type of objectives:

- Büchi objectives. Assume that each Player i wants to visit a set of states T_i infinitely often. In this case, R_i is the set of states s from which Player i has a strategy to enforce the objective (in LTL syntax) $\Box\Diamond T_i \ \vee \ \bigwedge_{j=1}^{j=n} \Diamond\Box\overline{T_j}$. This formula is equivalent to $\Box\Diamond T_i \ \vee \ \Diamond\Box \bigwedge_{j=1}^{j=n} \overline{T_j}$. This is equivalent to a disjunction of a Büchi and a co-Büchi objective, which is thus equivalent to a Streett objective with one Streett pair and can be solved in PTime, see e.g. [32].

- co-Büchi objectives. Assume that each Player i wants to eventually stay forever in a set of states T_i. In this case, R_i is the set of states s from which Player i has a strategy to enforce the objective (in LTL syntax) $\Diamond\Box T_i \ \vee \ \bigwedge_{j=1}^{j=n} \Box\Diamond\overline{T_j}$. The second part of the formula, i.e. $\bigwedge_{j=1}^{j=n} \Box\Diamond\overline{T_j}$ is a *generalized Büchi objective*. Using a standard construction, see e.g. Theorem 4.56 in [5], we can modify the game arena so that the generalized Büchi objective is transformed into a *classical* Büchi objective. This construction works as follows: it maintains a counter which counts modulo n and whose values determine the next set $\overline{T_j}$ to visit, once this set is visited, the counter is incremented. Now, the associated Büchi objective asks to visit infinitely often successive states in which the counter goes from value n to value 1. This ensures that each set $\overline{T_j}$, for $1 \leq j \leq n$, is visited infinitely often as requested. So, the transformation increases the size to the game arena by a factor of n and the objective that we get is, as in the previous case, a disjunction of a Büchi and a co-Büchi objective, which is thus equivalent to a Streett objective with one Streett pair and can be solved in PTime.

- For the parity case, the winning objectives for the retaliation sets can be encoded compactly as Muller objectives defined by a propositional formula using one proposition per state. Then they can be solved in PSpace using the algorithm of Emerson and Lei presented in [16].

□

Lower bounds. Let us now establish the lower bounds. The hardness for Büchi and co-Büchi objectives holds already for 2 players.

Lemma 6. *The problem of deciding the existence of a doomsday equilibrium in a 2-player game arena is* PTIME*-hard both for Büchi and co-Büchi winning objectives.*

PROOF. We explain the result for Büchi objectives (the proof for co-Büchi objectives is similar). To establish this result, we show how to reduce the problem of deciding the winner in a two-player zero-sum game with a Büchi objective (for Player 1), a PTIME-C problem [23], to the existence of a doomsday equilibrium in a two-player game arena with Büchi objectives. Let G be the two-player game, S its set of states, and T the set of states that Player 1 wants to visit infinitely often. We reduce the problem of deciding the existence of such a strategy to the existence of a doomsday equilibrium in the same game arena, where the objective of Player 1 is the original Büchi objective, i.e. $\mathsf{Büchi}(T)$, and the objective of Player 2 is trivial: $\mathsf{Büchi}(S)$. Clearly, as Player 2 will always satisfy his objective, Player 1 must have a winning strategy for $\mathsf{Büchi}(T)$ if a doomsday equilibrium exists (and vice versa) otherwise condition 2 would be violated. Indeed as Player 2 can never lose, Player 1 cannot retaliate and so he must have a winning strategy for his own objective, i.e. for $\mathsf{Büchi}(T)$. □

We now turn to the proof of lower bounds for parity objectives. We show below that we can reduce zero-sum two-player games with a conjunction of parity objectives (known to be CONP-HARD [13]) to the existence of a DE in a *three-player* game with parity objectives. Similarly, we can reduce the problem of deciding the winner in a *two-player* zero-sum game with a disjunction of parity objectives (known to be NP-HARD [13]) to the existence of a DE in a *two-player* game with parity objectives. The main idea in the two cases is to construct a game arena where one of the players can retaliate iff Player 1 in the original two-player zero-sum game has a winning strategy. The details of those reductions are given in the proofs of the following lemmas.

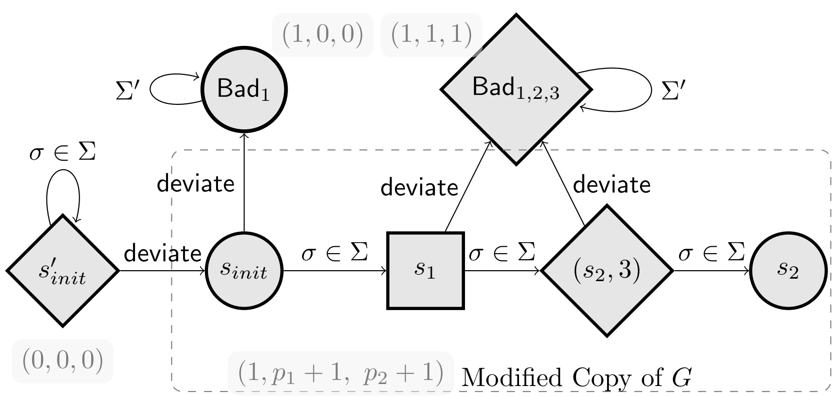

Figure 3: Structure of the reduction from generalized parity game with a conjunction of two parity objectives to the existence of a doomsday equilibrium with parity objectives.

Lemma 7 (CONP-HARDNESS). *The problem of deciding the existence of a doomsday equilibrium in a* 3*-player game arena with parity objectives is* CONP-HARD.

PROOF. Let $G = (S, \{S_A, S_B\}, s_{\mathsf{init}}, \Sigma, \Delta)$ be a two-player game and a conjunction of two parity objectives defined by the functions p_1 and p_2 that Player A wants to enforce, i.e. the objective of Player A is to ensure an outcome that satisfies the two parity objectives, while the objective of Player B is to ensure an outcome that violates at least one of the two parity objectives. W.l.o.g., we assume that $s_{\mathsf{init}} \in S_A$ and the turns of A and B alternate.

From G, we construct a 3-player game arena $G' = (S', \{S'_1, S'_2, S'_3\}, s'_{\mathsf{init}}, \Sigma', \Delta')$ (depicted in Fig. 3), with:

- the set of states $S' = \{s'_{\mathsf{init}}, \mathsf{Bad}_1, \mathsf{Bad}_{1,2,3}\} \cup S_A \cup S_B \cup (S_A \times \{3\})$, this set is partitioned as follows: $S_1 = S_A \cup \{\mathsf{Bad}_1\}$, $S_2 = S_B$, $S_3 = (S_A \times \{3\}) \cup \{s'_{\mathsf{init}}, \mathsf{Bad}_{1,2,3}\}$.

- the initial state is s'_{init},

- the alphabet of actions is $\Sigma' = \Sigma \cup \{\mathsf{deviate}\}$,

- and the transitions of the game G' are defined as follows:

  - For the state s'_{init}, for all $\sigma \in \Sigma$, $\Delta'(s'_{\mathsf{init}}, \sigma) = s'_{\mathsf{init}}$, and $\Delta'(s'_{\mathsf{init}}, \sigma) = s_{\mathsf{init}}$; i.e., the play stays in s'_{init}, unless Player 3 plays deviate in which case the play goes to s_{init} that is the copy of the initial state of the game arena G.
  - For all states $s \in S_A$, for all $\sigma \in \Sigma$, $\Delta'(s, \sigma) = \Delta(s, \sigma)$, and $\Delta'(s, \mathsf{deviate}) = \mathsf{Bad}_1$, so the transition function on the copy of G behaves from states owned by Player 1 as in the original game and it sends the game to Bad_1 if Player 1 plays the action deviate.
  - For all states $s \in S_B$, for all $\sigma \in \Sigma$, $\Delta'(s, \sigma) = (\Delta(s, \sigma), 3)$ and $\Delta'(s, \mathsf{deviate}) = \mathsf{Bad}_{1,2,3}$; i.e., if Player 2 plays an action from the game G, the effect is to send the game to the Player 3 copy of the same state as in the original game, if he deviates, the game reaches the sink state $\mathsf{Bad}_{1,2,3}$.
  - For all states $s \in S_A \times \{3\}$, for all $\sigma \in \Sigma$, $\Delta'((s, 3), \sigma) = s$ and $\Delta'((s, 3), \mathsf{deviate}) = \mathsf{Bad}_{1,2,3}$. So, if Player 3 plays an action $\sigma \in \Sigma$, he gives back the turn to Player 1, otherwise he sends the game to $\mathsf{Bad}_{1,2,3}$.
  - The states Bad_1 and $\mathsf{Bad}_{1,2,3}$ are absorbing.

- The parity functions $(p'_i)_{i=1,2,3}$ for the three players are defined to satisfy the following condition:

  - first, $p'_i(s'_{\mathsf{init}})$ is even for all $i = 1, 2, 3$ (so if the game stays there for ever, all the players satisfy their objectives).
  - second, in Bad_1 the parity functions return an even number for Player 2 and Player 3 but an odd number for Player 1, this ensures that Player 1 should never play the action deviate when the game is in the copy of G,
  - third, in $\mathsf{Bad}_{1,2,3}$ the parity functions are odd for all the Players. So whenever Player 2 and 3 play deviate all players loose,
  - finally, in the copy of G, the parity function is always odd for Player 1, and for all states $q \in S_A \cup S_B \cup (S_A \times \{3\})$, $p'_2(q) = p_1(s) + 1$ and $p'_3(q) = p_2(s) + 1$, where $s = q$ if $q \in S_A \cup S_B$, and s is such that $q = (s, 3)$ if $q \in S_A \times \{3\}$.

This concludes the reduction.

Clearly, since Player A and B always alternate their moves, in the copy of G, any play will eventually reach a state of Player 2 and a state of Player 3, so that they are always able to retaliate by playing the action deviate.

A doomsday equilibrium exists in G' iff Player 1 is also able to retaliate when the game enter the copy of G. But clearly, it is possible if and only if he has a strategy to ensure $\overline{\mathsf{parity}(p'_1)}$ and $\overline{\mathsf{parity}(p'_2)}$, or equivalently iff he has a strategy to ensure $\mathsf{parity}(p_1)$ and $\mathsf{parity}(p_2)$, iff Player A has a winning strategy in the game G for the conjunction of parity objectives p_1 and p_2. $\square$

Lemma 8 (NP-HARDNESS). *The problem of deciding the existence of a doomsday equilibrium in a 2-player game arena with parity objectives is* NP-HARD.

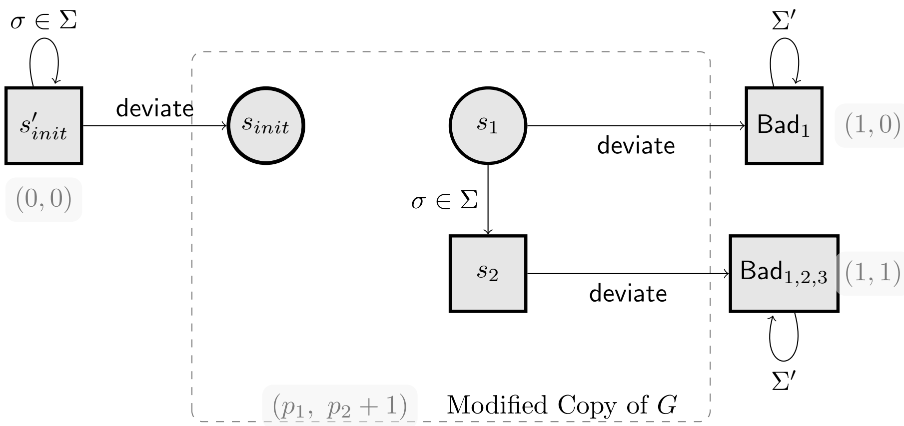

Figure 4: Structure of the reduction from generalized parity game with a disjunction of two parity objectives to doomsday equilibrium with parity objectives.

Proof. For this part, we need to show how to reduce the problem of deciding if Player A has a winning strategy in a two-player zero-sum game whose objective is defined by the disjunction of two parity objectives. The construction is based on the main ideas used in the coNP-hardness result given in the proof of lemma 7. Let $G = (S, s_{\mathsf{init}}, \Sigma, \Delta)$ be a two-player game with a *disjunction* of two parity objectives defined by the functions p_1 and p_2. The objective of Player A is to ensure an outcome that satisfies at least one of the two parity objectives (while the objective of Player B is to ensure an outcome that violates both parity objectives.)

From G, we construct a two-player game G' with parity objectives $(p'_i)_{i=1,2}$ (see Fig. 4). The game arena G' contains a copy of G plus three states s'_{init} (the initial state), Bad_1 and $\mathsf{Bad}_{1,2}$. The alphabet of actions is $\Sigma \cup \{\mathsf{deviate}\}$.

The partition of the state space is as follows: $S_1 = S_A$ and $S_2 = S_B \cup \{s'_{\mathsf{init}}\} \cup \{\mathsf{Bad}_1, \mathsf{Bad}_{1,2}\}$. The transitions are as follows: if Player 2 plays $\sigma \in \Sigma$ in s'_{init} then the game stays there, if he plays $\mathsf{deviate}$ then the game enters the copy of G. There the transition function for $\sigma \in \Sigma$ is defined as in G, and if Player 1 plays $\mathsf{deviate}$ then the game goes to Bad_1, and if Player 2 plays $\mathsf{deviate}$ then the game goes to $\mathsf{Bad}_{1,2}$.

The parity functions $(p'_i)_{i=1,2}$ are defined as follows: p'_1 returns an even number in s'_{init}, is equal to p_1 in the copy of G, returns an odd number in Bad_1 and $\mathsf{Bad}_{1,2}$. The function p'_2 returns an even number in s'_{init}, is equal to $p_2 + 1$ in the copy of G (so p'_2 is the complement of p_2), returns an even number in Bad_1 and an odd number in $\mathsf{Bad}_{1,2}$. This definition of p'_1 and p'_2 ensures that:

- the two players meet their parity objectives when the game always stays in s'_{init}.
- when the game enters the copy of G, Player 2 can always retaliate by playing $\mathsf{deviate}$, as the play then reaches $\mathsf{Bad}_{1,2}$ and the outcome is bad for both player.
- when the game enters the copy of G, Player 1 can retaliate, i.e. enforces $p'_1 \vee \overline{p'_2}$, which is equivalent to $p_1 \vee p_2$, if and only if Player A has a winning strategy for the disjunction of the parity objectives in the original game. Indeed, playing $\mathsf{deviate}$ is not an option for Player 1 in the copy of G, as then Bad_1 is reached and Player 2 wins while Player 1 loses.

So we conclude that there is a DE equilibrium in G' if and only if Player A has a winning strategy in the original game for $p_1 \vee p_2$. □

3.2. Reachability objectives

We now establish the complexity of deciding the existence of a doomsday equilibria in an n-player game with reachability objectives. We first establish an important property for reachability objectives:

Proposition 9. *Let $G = (S, \mathcal{P}, s_{\mathsf{init}}, \Sigma, \Delta)$ be a game arena, and $(T_i)_{1 \leq i \leq n}$ be n subsets of S. Let Λ be a doomsday equilibrium in G for the reachability objectives $(\mathsf{Reach}(T_i))_{1 \leq i \leq n}$. Let s be the first state in $\mathsf{outcome}(\Lambda)$ such that $s \in \bigcup_i T_i$. Then every player has a strategy from s, against all the other players, to reach his target set.*

Proof. Let $\Lambda = (\lambda_1, \lambda_2, \ldots, \lambda_n)$, and w.l.o.g. we can assume that $s \in T_1$. If some player, say Player 2, has no strategy from s to reach his target set T_2, then necessarily $s \notin T_2$ and by determinacy the other players have a strategy from s to make Player 2 lose. This contradicts the fact that Λ is a doomsday equilibrium as it means that λ_2 is not a retaliating strategy. $\square$

Lemma 10. *The problem of deciding the existence of a doomsday equilibrium in an n-player game with reachability objectives is in* PTime.

Proof. The algorithm consists in:

- First, computing the sets R_i from which player i can retaliate, i.e. the set of states s from which Player i has a strategy to force, against all other players, an outcome such that $\Diamond T_i \vee (\bigwedge_{j=1}^{j=n} \Box \overline{T_j})$. This set can be obtained by first computing the set of states $\langle\langle i \rangle\rangle \Diamond T_i$ from which Player i can force T_i to be reached. It is done in PTime by solving a classical two-player reachability game, see e.g. [23]. Then, one computes the set of states where Player i has a strategy λ_i such that $\mathsf{outcome}_i(\lambda_i) \models \Box((\bigcap_{j=1}^{j=n} \overline{T_j}) \vee \langle\langle i \rangle\rangle \Diamond T_i)$, that is to confine the plays in states that do not satisfy the reachability objectives of the adversaries or from where Player i can force its own reachability objective. Again this can be done in PTime by solving a classical two-player safety game.
- Second, checking the existence of some $i \in \{1, \ldots, n\}$ and some finite path π starting from s_{init} and that stays within $\bigcap_{j=1}^{j=n} R_j$ before reaching a state s such that $s \in T_i$ and $s \in \bigcap_{j=1}^{j=n} \langle\langle j \rangle\rangle \Diamond T_j$.

Let us now prove the correctness of our algorithm. From its output, we can construct the strategy profile $\Lambda = (\lambda_1, \lambda_2, \ldots, \lambda_n)$ where each λ_j ($j = 1, \ldots, n$) is as follows: follow π up to the point where either another player deviates and then play the retaliating strategy available in R_j, or to the point where some state $s \in T_i$ is visited, for some $i \neq j$, for the first time and then play according to a strategy (from s) that forces a visit to T_j no matter how the other players are playing. Clearly, Λ witnesses a DE. Indeed, if $s \in T_i$ is reached, then all players have a strategy to reach their target set (including Player i since $s \in T_i$). By playing so they will all eventually reach it. Before reaching s, if some of them deviates, the others have a strategy to retaliate as π stays in $\bigcap_{j=1}^{j=n} R_j$. The other direction follows from Proposition 9. $\square$

Lemma 11. *The problem of deciding the existence of a DE in a 2-player game with reachability objectives is* PTime-Hard.

Proof. The idea of the proof is similar to the one of Lemma 6. It is proved by a reduction from the And-Or graph reachability problem [23]. From an instance of the And-Or graph reachability problem, we construct a two-player game arena which is a copy of the And-Or graph. The first player owns the positions of the protagonist in the And-Or graph (i.e., the *Or* positions) and the second player owns the positions of the antagonist (i.e., the *And* positions). The objective of the first player coincides with the objective of the protagonist in the And-Or graph reachability problem, while the objective of the second player is trivial: his set of target states is the entire state space. So, clearly as the second player always wins for his objective, the only way to have a doomsday equilibrium in the constructed game is to have a winning strategy for the first player for his objective, which is equivalent to have a winning strategy for the protagonist in the And-Or graph reachability problem. $\square$

3.3. Safety Objectives

We establish the complexity of deciding the existence of a doomsday equilibrium in an n-player game with perfect information and safety objectives.

Lemma 12 (Easyness). *The existence of a doomsday equilibrium in an n-player game with safety objectives can be decided in* PSpace, *and in* PTime *for game arenas with a fixed number of players.*

Proof. We start with the general case where the number of players is not fixed and is part of the input. Let us consider an n-player game arena $G = (S, \mathcal{P}, s_{\mathsf{init}}, \Sigma, \Delta)$ and n safety objectives $\mathsf{safe}(T_1), \ldots, \mathsf{safe}(T_n)$ for $T_1 \subseteq S, \ldots, T_n \subseteq S$. The algorithm is composed of the following two steps:

- First, for each Player i, compute the set of states $s \in S$ in the game such that Player i can retaliate whenever necessary, i.e. the set of states s from where there exists a strategy λ_i for Player i such that $\mathsf{outcome}_i(\lambda_i)$ satisfies $\neg\Box T_i \rightarrow \bigwedge_{j=1}^{j=n} \neg\Box T_j$, or equivalently $\neg\Diamond\overline{T_i} \vee \bigwedge_{j=1}^{j=n} \Diamond\overline{T_j}$. This can be done in PSpace using a result by Alur et al. (Theorem 5.4 of [3]) on solving two-player games whose Player 1's objective is defined by Boolean combinations of LTL formulas that use only $\Diamond$ and $\wedge$. We denote by R_i the set of states in G where Player i has a strategy to retaliate.

  We note here that as R_i is the winning set for Player i for the objective $\neg\Diamond\overline{T_i} \vee \bigwedge_{j=1}^{j=n} \Diamond\overline{T_j}$ then it is a trap for the other players in the following precise sense: for all $s \in R_i$, either $s \in \bigcap_{j=1}^{j=n} \overline{T_j}$ or:

  1. if $s \in R_i \cap S_i$ then there exists $\sigma \in \Sigma$: $\Delta(s, \sigma) \in R_i$, i.e. Player i can stay within his retaliating region,
  2. if $s \notin R_i \cap S_i$, then for all $\sigma \in \Sigma$: $\Delta(s, \sigma) \in R_i$, i.e. the other players cannot escape from the retaliating region R_i.

- Second, verify whether there exists an infinite path ρ in $\bigcap_{i=1}^{i=n}(\mathsf{safe}(T_i) \cap R_i)$.

Now, let us establish the correctness of this algorithm. Assume that an infinite path ρ exists in $\bigcap_{i=1}^{i=n}(\mathsf{safe}(T_i) \cap R_i)$. The strategies λ_i for each Player i are defined as follows: play the moves that are prescribed by ρ as long as every other players do so, and as soon as the play deviates from ρ, play the retaliating strategy.

It is easy to see that the profile of strategies $\Lambda = (\lambda_1, \lambda_2, \ldots, \lambda_n)$ is a DE. Indeed, the states are all safe for all players as long as they play their strategies. Moreover, assume that some player deviates from ρ, and let $s \in S_j$, for some j, be the state of ρ where the first deviation occurs, i.e. Player j is the player who deviates from ρ. Since the play ρ is within $\bigcap_{i=1}^{i=n} R_i$, by the trap property defined above, we know that the state that is reached after s is still in $\bigcap_{i=1, i\neq j}^{i=n} R_i$ and therefore the other players can retaliate.

Second, assume that $\Lambda = (\lambda_1, \lambda_2, \ldots, \lambda_n)$ is a DE in the n-player game G for the safety objectives $(\mathsf{safe}(T_i))_{1\leq i\leq n}$. Let $\rho = \mathsf{outcome}(\lambda_1, \lambda_2, \ldots, \lambda_n)$. By definition of doomsday equilibrium, we know that all states appearing in ρ satisfy all the safety objectives, i.e. $\rho \models \bigwedge_{i=1}^{i=n} \Box T_i$. Let us show that the play also remains within $\bigcap_{i=1}^{i=n} R_i$. Let s be a state of ρ, $i \in \{1, \ldots, n\}$, and π the finite prefix of ρ up to s. By definition of DE we have $\mathsf{outcome}(\lambda_i) \models \Box T_i \vee \bigwedge_{j=1}^{j=n} \Diamond\overline{T_j}$. Therefore for all outcomes ρ' of λ_i in G_s, $\pi\rho' \models \Box T_i \vee \bigwedge_{j=1}^{j=n} \Diamond\overline{T_j}$. Moreover, $\pi \models \bigwedge_{j=1}^{j=n} \Box T_j$ since it is a prefix of ρ. Therefore $\rho' \models \Box T_i \vee \bigwedge_{j=1}^{j=n} \Diamond\overline{T_j}$ and $s \in R_i$. Since it holds for all $i \in \{1, \ldots, n\}$, we get $s \in \bigcap_{i=1}^{i=n} R_i$.

Let us now turn to the case where the number of players is fixed. Then clearly, in the construction above, all the LTL formulas are of fixed size and so all the associated games can then be solved in polynomial time. $\Box$

For the general case, we present a reduction from the problem of deciding the winner in a zero-sum two-player game with a conjunction of k reachability objectives (aka generalized reachability games), which is a PSpace-C problem [4]. The idea of the reduction is to construct a non-zero sum $(k+1)$-player game where one of the players has a retaliating strategy iff there is a winning strategy in the generalized reachability game. When the number of players is fixed, PTime-Hardness is proved by an easy reduction from the And-Or graph reachability problem [23].

Lemma 13 (Hardness). *The problem of deciding the existence of a doomsday equilibrium in an n-player game with safety objectives is* PSpace-Hard, *and* PTime-Hard *when the number of players is fixed.*

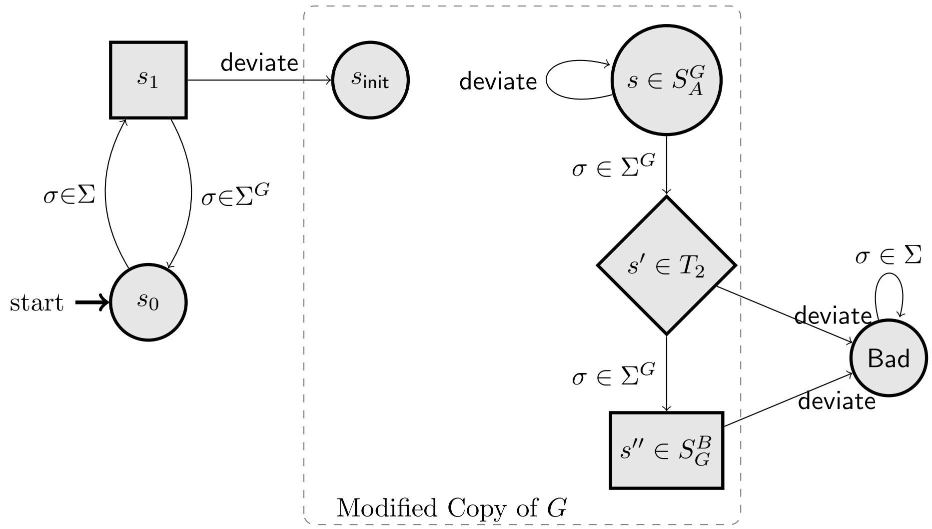

Figure 5: Structure of the reduction from multi-reachability game to doomsday equilibrium with safety objectives. Round nodes denote Player 0's states and rectangular nodes denote Player 1's states.

PROOF. We reduce the two-player *multi-reachability* problem to our problem, PSPACE-HARDNESS follows. Let $G = (S^G, \{S_A^G, S_B^G\}, s_{\mathsf{init}}^G, \Sigma^G, \Delta^G)$ be a two-player (Player A and Player B) game arena. Let $\mathcal{T} = \{T_1, T_2, \ldots, T_k\}$ be a family of subsets of S^G supposed to be pairwise disjoint (w.l.o.g.). Also wlog we assume that $T_i \subseteq S_B^G$ for all $i \in \{1, \ldots, k\}$. In a multi-reachability game, the objective of Player A is to visit each T in $\mathcal{T}$, while Player B tries to avoid at least one of the subsets in $\mathcal{T}$. So, multi-reachability games are two-player zero sum games where the winning plays for Player A are

$$\{\rho = s_0 s_1 \ldots s_n \cdots \in \mathsf{Plays}(G) \mid \forall i \cdot 1 \leq i \leq k \cdot \exists j \geq 0 \cdot s_j \in T_i\}.$$

It has been shown that the multi-reachability problem for two-player games is PSPACE-C [4, 17].

From $G = (S^G, \{S_A^G, S_B^G\}, s_{\mathsf{init}}^G, \Sigma^G, \Delta^G)$ and $\mathcal{T} = \{T_1, T_2, \ldots, T_k\}$ that define a multi-reachability game, we construct a game arena $G' = (S, \{S_0, S_1, \ldots, S_k\}$,
$s_0, \Sigma, \Delta)$ with $k+1$ players and a set of $k+1$ safety objectives $\mathsf{Safe}_0, \ldots, \mathsf{Safe}_k$ such that Player A wins the multi-reachability objective defined by G and $\mathcal{T}$ iff there exists a doomsday equilibrium in G' for the safety objectives $\mathsf{Safe}_0, \ldots, \mathsf{Safe}_k$.

The structure of the reduction is depicted in Fig. 5. The state space of G' is composed of three parts: an initial part on the left, a modified copy of G, and a part on the right. The set S of states is $\{s_0, s_1\} \cup S_A^G \cup S_B^G \cup \{\mathsf{Bad}\}$. This set of states is partitioned as follows: $S_0 = \{s_0\} \cup S_A^G \cup \{\mathsf{Bad}\}$, $S_1 = S_B^G \setminus \bigcup_{i=2}^{i=k} T_i$ and for all i, $2 \leq i \leq k$, $S_i = T_i$.
The sets of safety objectives are defined as follows: $\mathsf{Safe}_0 = \mathsf{safe}(\{s_0, s_1\})$, and for all $i \in \{1, 2, \ldots, k\}$, $\mathsf{Safe}_i = \mathsf{safe}(S \setminus (\{\mathsf{Bad}\} \cup T_i))$. The alphabet of actions is $\Sigma = \Sigma^G \cup \{\mathsf{deviate}\}$, and the transition function is defined as follows:

- $\Delta(s_0, \sigma) = s_1$, for all $\sigma \in \Sigma$,
- $\Delta(s_1, \sigma) = \begin{cases} s_0 & \text{if } \sigma \in \Sigma^G \\ s_{\mathsf{init}}^G & \text{if } \sigma = \mathsf{deviate} \end{cases}$
- for all $s \in S_A^G \cup S_B^G$:
  - for all $\sigma \in \Sigma^G$, $\Delta(s, \sigma) = \Delta^G(s, \sigma)$
  - for the letter $\mathsf{deviate}$: for all states $s \in S_A^G$, $\Delta(s, \mathsf{deviate}) = s$, and for all $s \in S_B^G$, $\Delta(s, \mathsf{deviate}) = \mathsf{Bad}$

- Bad is a sink state.

Now, let us justify this construction. First, assume that, in the two-player game arena $G = (S^G, \{S_1^G, S_2^G\}, s_{\mathsf{init}}^G, \Sigma^G, \Delta^G)$ with the multi-reachability objective given by $\mathcal{T} = \{T_1, T_2, \dots, T_k\}$, Player A has a winning strategy. In that case, we show that there exists a doomsday equilibrium in the game G' for the safety objectives $(\mathsf{Safe}_i)_{0 \leq i \leq k}$. To establish the existence of a doomsday equilibrium, we consider the strategy profile $\Lambda = (\lambda_0, \lambda_1, \dots, \lambda_k)$ whose strategies respect the following conditions:

- If all the players follows the strategy profile Λ, the outcome of the game is $(s_0 \cdot s_1)^\omega$, i.e. Player 1 avoids to play $\mathsf{deviate}$ in s_1.

- Whenever player 1 plays $\mathsf{deviate}$ in s_1, then the game enters the sub game of G' corresponding to G, and the game thus enters an unsafe state for Player 0 (as s_{init} is not part of Safe_0). From there, Player 0 must retaliate by forcing a visit to each set in $\mathcal{T} = \{T_1, T_2, \dots, T_k\}$ to make sure that all the other players lose. By hypothesis, in G, Player A has a winning strategy for the multi-reachability objective, so we know that if the other players play letters that are in Σ^G then all sets in $\mathcal{T}$ will eventually be visited when Player 0 plays according to the winning strategy of Player A in G. On the other hand, if the letter $\mathsf{deviate}$ is played then the game goes to the state Bad where all the safety objectives are violated. So, we have established that Player 0 can retaliate if he plays as Player A in the copy of G. Now, let us consider all the other players. According to the definition of the transition function, Player i has the option to retaliate whenever he enters its unsafe set T_i by choosing the action $\mathsf{deviate}$ and so force a visit to Bad. So, all other players have also the ability to retaliate whenever they enter their unsafe region.

So, we have established that $(\lambda_0, \lambda_1, \dots, \lambda_k)$ witnesses a doomsday equilibrium in G'.

Now, let us consider the other direction. Let $(\lambda_0, \lambda_1, \dots, \lambda_k)$ be a profile of strategies which witnesses a doomsday equilibrium for G' and the safety objectives given by the subsets of plays $(\mathsf{Safe}_i)_{i=0,..,k}$. In that case, if we consider a prefix of play that enters for the first time the state s_{init}, we know by definition of doomsday equilibrium that Player 0 has a strategy to retaliate against any strategies of the adversaries. If all the other players choose their letters in Σ^G then it should be the case that the play visits all the sets in $\mathcal{T}$. So, this clearly means that Player A has a winning strategy in G for the multi-reachability objective defined by $\mathcal{T}$, this strategy simply follows the strategy λ_0 in the copy of G. $\square$

3.4. LTL objectives

In this section, we show that the problem of deciding the existence of a DE for LTL objectives is 2EXPTIME-C. As a first solution, we define a reduction to parity objectives. For this reduction, we use DPW automata constructed from LTL formulas, so we need to apply the Safra determinization construction [36] or a variant. From this solution, we can bound the memory that is needed by the players to play a DE. Using that bound on memory, we provide a second solution in the form of a Safraless procedure, in the spirit of [26], that is based on a reduction to safety objectives as in [20, 19].

Recduction to parity objectives. Given a labelled n-player arena $G = (S, \mathcal{P}, s_{\mathsf{init}}, \Sigma, \Delta, \mathbb{P}, \mathcal{L})$ and LTL objectives $(\varphi_i)_{1 \leq i \leq n}$ over $\mathbb{P}$, we construct another n-player arena $G' = (S', s'_{\mathsf{init}}, \Sigma', \Delta')$ with *parity* objectives $(\varphi'_i)_{1 \leq i \leq n}$ such that there exists a DE in G if and only if there is one in G'. To construct G', we first construct a series of (total) deterministic parity automata from the LTL formulas $(\varphi_i)_{1 \leq i \leq n}$: $\mathcal{A}_{\mathsf{i}} = (Q^{\mathsf{i}}, q^{\mathsf{i}}_{\mathsf{init}}, 2^{\mathbb{P}}, \delta^{\mathsf{i}}, p^{\mathsf{i}})$ is such that $L(\mathcal{A}_{\mathsf{i}}) = [\![\varphi_i]\!]$, i.e. $\mathcal{A}_{\mathsf{i}}$ accepts exactly all the sequences of labels in which either player i wins. Then, we define the game arena $G' = (S', s'_{\mathsf{init}}, \Sigma', \Delta')$ as the product of G with those automata. The elements of G' are as follows:

- the set of states $S' = S \times Q^1 \times Q^2 \times \dots Q^{\mathsf{n}}$, i.e. the set of states of G' is the cartesian product of the set of states of G and the state spaces of all the automata for the LTL objectives;

- the initial state is $s'_{\mathsf{init}} = (s_{\mathsf{init}}, q^1_{\mathsf{init}}, q^2_{\mathsf{init}}, \dots, q^{\mathsf{n}}_{\mathsf{init}})$;

- the set of actions of S' is $\Sigma' = \Sigma$, so it is unchanged;
- the transition relation Δ' is defined as follows: for all $(s, q^1, q^2, \ldots, q^\mathsf{n}) \in S'$, $\sigma \in \Sigma$,

$$\Delta'((s, q^1, q^2, \ldots, q^\mathsf{n}), \sigma)$$
$$= (\Delta(s, \sigma), \delta^1(q^1, \mathcal{L}(s)), \delta^2(q^2, \mathcal{L}(s)), \ldots, \delta^\mathsf{n}(q^\mathsf{n}, \mathcal{L}(s)))$$

- the parity functions φ'_i are as follows: for all states $s' = (s, q^1, \ldots, q^\mathsf{n})$, $\varphi'_i(s') = p^\mathsf{i}(q^\mathsf{i})$.

Clearly, a play in G' is winning for φ'_i iff its projection in G corresponds exactly to a play that is winning for the LTL objective φ_i. As a direct consequence, we have that there is a correspondence between the DE in G and G'.

Lemma 14. *There exists a DE in the n-player labelled game arena $G = (S, s_{\mathsf{init}}, \Sigma, \Delta, \mathbb{P}, \mathcal{L})$ with LTL objectives $(\varphi_i)_{1 \le i \le n}$ if and only if there exists a DE in the n-player arena $G' = (S', s'_{\mathsf{init}}, \Sigma', \Delta')$ with* parity *objectives $(\varphi'_i)_{1 \le i \le n}$.*

This reduction from LTL objectives to parity objectives gives us a 2ExpTime procedure and it is worst-case optimal.

Theorem 15. *The problem of deciding the existence of a doomsday equilibrium in a labelled n-player game arena is* 2ExpTime-C *for LTL objectives.*

Proof. Lemma 14 justifies the correctness of the construction. We need to make precise the size of G' and the number of colors that are used in the parity conditions in order to give precise upper bounds for the complexity of our reduction. First, the size of all the automata constructed for the LTL formulas are bounded doubly exponentially in the size of the LTL formulas $(\varphi_i)_{1 \le i \le n}$, and each of them uses a number of colors for their parity condition which is bounded exponentially in the size of those formulas, those bounds on the translation from LTL to deterministic parity automata can be found e.g. in [26]. Let B_A be the size of the largest automaton, and B_d be the largest number of colors. Clearly, from the construction of G', its state space is at most $\mathbf{O}(|S| \cdot {B_A}^n)$ and each of the parity objectives uses at most B_d colors.

In the proof of Lemma 5, we have shown that to search for a DE for parity objectives, we need to compute the retaliation sets of each player. This can be done with the Emerson and Lei algorithm [16]. As the state space is linear in the state space of the game arena, doubly exponential in the size of the LTL objectives, and the parity conditions contains at most exponentially many colors in the size of the LTL formulas, then we know that the overall complexity of computing these is bounded by a double exponential in the size of the original problem. This gives the upper bound.

The lower bound is easily obtained by a reduction from the realizability problem for LTL, which is known to be 2ExpTime-Hard [33]. The proof closely follows the proof idea of lemma 6 and lemma 11. □

Bounding the memory in strategies. We now refine the complexity analysis of Theorem 15, and give a bound on the size of the memory of strategies that form a DE for LTL objectives. The following lemma shows that there always exists DE with bounded memory for parity objectives.

Lemma 16. *If there is a DE equilibrium in n-player arena $G = (S, \mathcal{P}, s_{\mathsf{init}}, \Sigma, \Delta)$ with the parity objectives $(\varphi_i)_{1 \le i \le n}$, then there is a profile $(\lambda_1, \lambda_2, \ldots, \lambda_n)$ where the memory used by each strategy λ_i is bounded polynomially in the number of states in G, exponentially in the number of players, and exponentially in the number of colors used in the parity conditions.*

Proof. Parity objectives are tail objectives, so to compute a strategy profile that forms a DE, we can apply the algorithm described in the proof of Lemma 5.

That algorithm first computes the set R_i for each player together with a strategy for retaliation. Then the algorithm looks for a path that is winning for every players within the intersection of the sets R_i. If such path exists, then the DE equilibrium is composed of strategies that follow this winning path for ever, or up to the point one of the player deviates, and in this case all the other players switch to their own retaliating strategy. So the memory needed for each player is thus bounded by the size of the winning path plus the size of the retaliating strategy. We now bound those two sizes.

For the size of the retaliation, we observe that the objective for player i is $\varphi_i \vee \bigwedge_{j\neq i,1\leq j\leq n} \varphi_j$, which is a Boolean combination of parity conditions. This Boolean combination can be rewritten as $\bigwedge_{j\neq i,1\leq j\leq n}(\varphi_i \vee \varphi_j)$. Now, each conjunct $\varphi_i \vee \varphi_j$ can be seen as a Rabin objective with $\mathbf{O}(d)$ pairs (where d is the maximal color index used in the two parity functions). This Rabin condition can be turned in to a Streett condition in a game with $\mathbf{O}(m \cdot 2^{d \log d})$ states and $\mathbf{O}(d)$ pairs using the *Index Appearance Record* construction provided in [36]. So, to compute the retaliating region, we need to solve a Streett game with $\mathbf{O}(m \cdot 2^{d \log d})$ states and $\mathbf{O}(d)$ pairs. Optimal strategies in Streett games can be implemented using finite-memory, and the memory size is bounded by $\mathbf{O}(d!)$, and those strategies can be constructed in time exponential in the number of pairs and polynomial in the number of states in the game. So, it means that for our games, we obtain that the region R_i can be computed in time polynomial in m (the state space of the arena) and exponential in d (the largest color index used in the winning condition for each player in the original game). So the memory needed for the retaliating strategy is $\mathbf{O}(m \cdot 2^{d \log d} \cdot d!)$.

Once we have computed, the retaliating regions, we search for a path in the intersection of all those regions that pleases all the players in G. This can be done by considering the subgraph of the arena G defined by $\bigcap_{1\leq i\leq n} R_i$. Looking for an infinite path which is winning for all players can be seen as looking for a accepting run in a deterministic Streett automaton. Indeed, each parity condition ϕ_i can be translated into a Streett condition with $\mathbf{O}(d)$ pairs and the union of those sets of pairs is equivalent to the conjunction of the parity conditions. It is known that the non-emptiness of a Streett automaton with α states and β pairs is witnessed by a lasso path uv^ω such that $|u| + |v| \leq \alpha\beta^2\beta!$ (see e.g. [31]).

So the overall memory requirement for each player i is bounded by the size of the retaliating strategy which is $\mathbf{O}(m \cdot 2^{d \log d} \cdot d!)$ and by the size of the outcome of the DE which is bounded by $\mathbf{O}(m \cdot (n \cdot d)^2 (n \cdot d)!)$.

□

From the reduction to parity objectives given in the proof of Lemma 14 and the previous lemma, we obtain the following corollary.

Corollary 17. *If a labelled n-player arena $G = (S, \mathcal{P}, s_{\mathsf{init}}, \Sigma, \Delta, \mathbb{P}, \mathcal{L})$ with LTL objectives $(\varphi_i)_{1\leq i\leq n}$ over $\mathbb{P}$ has a DE equilibrium then there exists a profile of strategies $(\lambda_1, \lambda_2, \ldots, \lambda_n)$ which is a DE and each strategy λ_i can be encoded as a Moore machine of size which his at most linear in the size of G, doubly exponential in the size of the largest formula in $(\varphi_i)_{1\leq i\leq n}$, and exponential in the number of players n.*

Safraless procedures - main ideas. Corollary 17 gives a bound on the memory size of strategies in a DE for LTL objectives. We exploit this bound to define a procedure that avoids the construction of DPW, and so avoids the use of the Safra determinization construction which is notoriously difficult to implement efficiently [1]. Our Safraless procedure is based on universal co-Büchi automata as suggested in [26] and leads to solving safety games as in [20, 18]. The idea underlying the safety game is as follows.

The central idea of Safraless methods, as first proposed in [26], is to exploit the fact that when a player wins an ω-regular objective in a two-player game G, such as an LTL objective φ, then he has a *finite-memory* winning strategy to enforce this objective[2]. The existence of a *finite-memory* winning strategy allows us to strengthen the LTL objective into a simpler objective, here a *safety objective* as in [20, 18]. The justification of this reduction follows the following line of arguments.

First, given an LTL objective φ, we can construct a *universal* co-Büchi automaton A_φ that accepts all the plays that satisfy φ, see e.g. [5]. This automaton has a size which is at most exponential in the size of the LTL formula.

[2]For LTL objective, it is known that a memory which is doubly exponential in the size of the LTL formula and linear in the size of the game structure suffices.

Second, assume that λ is a finite-memory winning strategy for the objective φ. Assume that λ uses a memory of size m, that A_φ has M_A states, and that the game G has M_G states. Then if we take the synchronised product of the strategy λ with the universal co-Büchi automaton A_φ, and with the game structure G, we obtain a graph with $m \cdot M_A \cdot M_G$ states. As λ is winning for the objective defined by the universal co-Büchi automaton A_φ, we know that all reachable cycles of this graph are free of accepting states in the automaton component (as on all runs of A_φ only finitely many accepting states should be visited on outcomes compatible with λ). But if no reachable cycles contain accepting states, the maximal number of accepting states seen on a play compatible with λ is bounded by the size of the graph, that is by $m \cdot M_A \cdot M_G$. So, λ ensures that the number of accepting states visited on a run that monitors a play compatible with λ is bounded by $m \cdot M_A \cdot M_G$. As a consequence, we can replace the co-Büchi objective (i.e. visits a finite number of times the accepting states) by a stronger safety objective which asks to visit accepting states at most $m \cdot M_A \cdot M_G$ times in A_φ). The strategy λ is indeed enforcing this stronger objective.

Formal construction. As our procedure follows closely the previous works mentioned above, here we only sketch the main ideas underlying the correctness and completeness arguments.

First, we need the following notation. Given a UcoBW $\mathcal{B}$, remember that $\mathcal{L}(\mathcal{B})$ denotes its language. Given a value K, we note $\mathcal{L}_K(\mathcal{B})$ the language defined by $\mathcal{B}$ when it is interpreted as a $\mathsf{U}K\mathsf{coBW}$, for which an accepting runs is allowed to visit at most K times a Büchi state (instead of a finite number of time as in the classical co-Büchi acceptance condition). Clearly we have that $\mathcal{L}_K(\mathcal{B}) \subseteq \mathcal{L}(\mathcal{B})$.

Given a labelled n-player arena $G = (S, \mathcal{P}, s_{\mathsf{init}}, \Sigma, \Delta, \mathbb{P}, \mathcal{L})$ and LTL objectives $(\varphi_i)_{1 \leq i \leq n}$ over $\mathbb{P}$, the algorithm follows the following steps:

1. We first construct the universal co-Büchi automata for each of the following formulas:
   - $\bigwedge_{i=1}^{i=n} \varphi_i$, and denote it by $\mathcal{B}^{\mathsf{all}}$;
   - for all i, $1 \leq i \leq n$, $\varphi_i \vee \bigwedge_{j \neq i, 1 \leq j \leq n} \neg\varphi_j$, and denote it by $\mathcal{B}^{\mathsf{r_i}}$.
2. We fix a positive integer K
3. Each automaton $\mathcal{B}^{\mathsf{r_i}}$ together with value K, defines a safety language and thus a safety game when used as an observer for G. For each such safety game, we compute Λ_i which is the most general strategy[3] for player i in G to enforce the safety language defined by $\mathcal{B}^{\mathsf{r_i}}$ and K.
4. We check if there is a profile of strategies $(\lambda_1, \lambda_2, \dots, \lambda_n)$ such that for all i, $1 \leq i \leq n$, $\lambda_i \in \Lambda_i$, and such that the outcome of $(\lambda_1, \lambda_2, \dots, \lambda_n)$ in G is satisfying $\bigwedge_{i=1}^{i=n} \varphi_i$, i.e. it is accepted by $\mathcal{B}^{\mathsf{all}}$, return **Yes** if such profile exists and No otherwise.

Clearly, all the steps above avoid Safra's determinization. We now show that positive answers of the procedure are correct for any value of K, and negative answers are correct for sufficiently large values of K.

Theorem 18. *When the Safraless procedure above returns* Yes *then there exists a DE for the labelled n-player arena $G = (S, s_{\mathsf{init}}, \Sigma, \Delta, \mathbb{P}, \mathcal{P})$ with LTL objectives $(\varphi_i)_{1 \leq i \leq n}$, and there exists a K which is at most linear in the size of G, doubly exponential in the size of the LTL objectives, and exponential in the number of players, such that if the procedure returns* No *for K, then there is no DE equilibrium in G for the LTL objectives $(\varphi_i)_{1 \leq i \leq n}$.*

Proof. Clearly, if the procedure returns Yes then there is a profile of strategies $(\lambda_1, \lambda_2, \dots, \lambda_n)$ such that all the players win their objective when this profile is played and furthermore every λ_i enforce the language $\mathcal{L}_K(\mathcal{B}^{\mathsf{r_i}})$, and we know that $\mathcal{L}_K(\mathcal{B}^{\mathsf{r_i}}) \subseteq \mathcal{L}(\mathcal{B}^{\mathsf{r_i}}) = [\![\varphi_i \bigwedge_{j \neq i, 1 \leq j \leq n} \neg\varphi_j]\!]$. So, $(\lambda_1, \lambda_2, \dots, \lambda_n)$ is a DE equilibrium for the LTL objectives $(\varphi_i)_{1 \leq i \leq n}$.

Now, if there exists a DE, by Corollary 17, we know that there exists a DE with a profile of strategies $(\lambda_1, \lambda_2, \dots, \lambda_n)$ where each λ_i has a memory size bounded linearly in the size of G, doubly exponentially in the size of the LTL objectives, and exponentially in the number of players. Let B denotes that bound.

[3]In a safety game, the set of all winning strategies is define by the maximal subset of edges that are safe in the positions of the protagonist. So the notion of most general strategy is well defined.

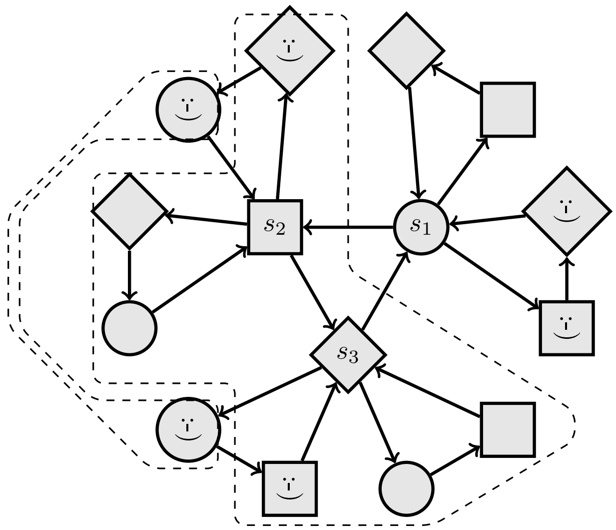

Figure 6: Game arena with imperfect information and Büchi objectives. Only indistinguishable states of Player 1 (circle) are depicted. Observations are symmetric for the other players.

The outcome of the profile $(\lambda_1, \lambda_2, \ldots, \lambda_n)$ satisfies the objective of each player and each strategy λ_i enforces $\mathcal{L}(\mathcal{B}^{\mathsf{r_i}}) = [\![\varphi_i \bigwedge_{j\neq i, 1\leq j\leq n} \neg\varphi_j]\!]$, but as λ_i has memory size at most B, then for $K_i = B_i \cdot |G|$, we have that strategy λ_i also enforces $\mathcal{L}_{K_i}(\mathcal{B}^{\mathsf{r_i}})$. So, if we take K as the largest of those K_i and the procedure returns No then we know that there is no DE in G for the LTL objectives $(\varphi_i)_{1\leq i\leq n}$. □

4. Complexity of DE for Imperfect Information Games

In this section, we define n-player game arenas with imperfect information. We adapt to this context the notions of observation, observation of a play, observation-based strategies, and we study the notion of doomsday equilibria when players are restricted to play observation-based strategies.

Game arena with imperfect information. An n-player game arena with *imperfect information* is a tuple $G = (S, \mathcal{P}, s_{\mathsf{init}}, \Sigma, \Delta, (O_i)_{1\leq i\leq n})$ such that $(S, \mathcal{P}, s_{\mathsf{init}}, \Sigma, \Delta)$ is a game arena (of perfect information) and for all i, $1 \leq i \leq n$, $O_i \subseteq 2^S$ is a *partition* of S. Each block in O_i is called an *observation* of Player i. We assume that the players play in a predefined order[4]: for all $i \in \{1, \ldots, n\}$, all $q \in S_i$ and all $\sigma \in \Sigma$, $\Delta(q, \sigma) \in S_{(i \bmod n)+1}$.

Observations. For all $i \in \{1, \ldots, n\}$, we denote by $O_i(s) \subseteq S$ the block in O_i that contains s, that is the observation that Player i has when he is in state s. We say that two states s, s' are *indistinguishable* for Player i if $O_i(s) = O_i(s')$. This defines an equivalence relation on states that we denote by $\sim_i$. The notions of plays and prefixes of plays are slight variations from the perfect information setting: a play in G is a sequence $\rho = s_0, \sigma_0, s_1, \sigma_1, \dots \in (S \cdot \Sigma)^\omega$ such that $s_0 = s_{\mathsf{init}}$, and for all $j \geq 0$, we have $s_{j+1} = \Delta(s_j, \sigma_j)$. A prefix of play is a sequence $\pi = s_0, \sigma_0, s_1, \sigma_1, \ldots, s_k \in (S \cdot \Sigma)^* \cdot S$ that can be extended into a play. As in the perfect information setting, we use the notations $\mathsf{Plays}(G)$ and $\mathsf{PrefPlays}(G)$ to denote the set of plays in G and its set of prefixes, and $\mathsf{PrefPlays}_i(G)$ for the set of prefixes that end in a state that belongs to Player i. While actions are introduced explicitly in our notion of play and prefix of play, their visibility is limited by the notion of observation. The *observation* of a play $\rho = s_0, \sigma_0, s_1, \sigma_1, \ldots$ by Player i is the

[4]This restriction is not necessary to obtain the results presented in this section (e.g. Theorem 23) but it makes some of our notations lighter.

infinite sequence written $\mathsf{Obs}_i(\rho) \in (O_i \times (\Sigma \cup \{\tau\}))^\omega$ such that for all $j \geq 0$, $\mathsf{Obs}_i(\rho)(j) = (O_i(s_j), \tau)$ if $s_j \notin S_i$, and $\mathsf{Obs}_i(\rho)(j) = (O_i(s_j), \sigma_j)$ if $s_j \in S_i$. Thus, only actions played by Player i are visible along the play, and the actions played by the other players are replaced by τ. The observation $\mathsf{Obs}_i(\pi)$ of a prefix π is defined similarly. Given an infinite sequence of observations $\eta \in (O_i \times (\Sigma \cup \{\tau\}))^\omega$ for Player i, we denote by $\gamma_i(\eta)$ the set of plays in G that are compatible with η, i.e. $\gamma_i(\eta) = \{\rho \in \mathsf{Plays}(G) \mid \mathsf{Obs}_i(\rho) = \eta\}$. The functions γ_i are extended to prefixes of sequences of observations naturally.

Observation-based strategies and doomsday equilibria. A strategy λ_i of Player i is *observation-based* if for all prefixes of plays $\pi_1, \pi_2 \in \mathsf{PrefPlays}_i(G)$ such that $\mathsf{Obs}_i(\pi_1) = \mathsf{Obs}_i(\pi_2)$, it holds that $\lambda_i(\pi_1) = \lambda_i(\pi_2)$, i.e. while playing with an observation-based strategy, Player i plays the same action after indistinguishable prefixes. A strategy profile Λ is observation-based if each Λ_i is observation-based. Winning objectives, strategy outcomes and winning strategies are defined as in the perfect information setting. We also define the notion of outcome relative to a prefix of a play. Given an observation-based strategy λ_i for Player i, and a prefix $\pi = s_0, \sigma_0, \dots, s_k \in \mathsf{PrefPlays}_i(G)$, the strategy λ_i^π is defined for all prefixes $\pi' \in \mathsf{PrefPlays}_i(G_{s_k})$ where G_{s_k} is the game arena G with initial state s_k, by $\lambda_i^\pi(\pi') = \lambda_i(\pi \cdot \pi')$. The set of outcomes of the strategy λ_i *relative to* π is defined by $\mathsf{outcome}_i(\pi, \lambda_i) = \pi \cdot \mathsf{outcome}_i(\lambda_i^\pi)$.

The notion of doomsday equilibrium is defined as for games with perfect information but with the additional requirements that *only* observation-based strategies can be used by the players. Given an n-player game arena with imperfect information G and n winning objectives $(\varphi_i)_{1\leq i\leq n}$ (defined as in the perfect information setting), we want to solve the problem of deciding the existence of an *observation-based strategy profile* Λ which is a doomsday equilibrium in G for $(\varphi_i)_{1\leq i\leq n}$.

Example 19. Fig. 6 depicts a variant of the example in the perfect information setting, with imperfect information. In this example let us describe the situation for Player 1. It is symmetric for the other players. Assume that when Player 2 or Player 3 send their information to Player 1 (modeled by a visit to his happy states), Player 1 cannot distinguish which of Player 2 or 3 has sent the information, e.g. because of the usage of a cryptographic primitive. Nevertheless, let us show that there exists doomsday equilibrium. Assume that the three players agree on the following protocol: Player 1 and 2 send their information but not Player 3.

Let us show that this sequence witnesses a doomsday equilibrium and argue that this is the case for Player 1. From the point of view of Player 1, if all players follow this profile of strategies then the outcome is winning for Player 1. Now, let us consider two types of deviation. First, assume that Player 2 does not send his information (i.e. does not visit the happy states). In that case Player 1 will observe the deviation and can retaliate by not sending his own information. Therefore all the players are losing. Second, assume that Player 2 does not send his information but Player 3 does. In this case it is easy to verify that Player 1 cannot observe the deviation and so according to his strategy will continue to send his information. This is not problematic because all the plays that are compatible with Player 1's observations are such that: (i) they are winning for Player 1 (note that it would be also acceptable that all the sequence are either winning for Player 1 or losing for all the other players), and (ii) Player 1 is always in position to retaliate along this sequence of observations. In our solution below these two properties are central and will be called *doomsday compatible* and *good for retaliation*. □

Generic Algorithm. We present a generic algorithm to test the existence of an observation-based doomsday equilibrium in a game of imperfect information. To present this solution, we need two additional notions: sequences of observations which are *doomsday compatible* and prefixes which are *good for retaliation*. These two notions are defined as follows. In a game arena $G = (S, \mathcal{P}, s_{\mathsf{init}}, \Sigma, \Delta, (O_i)_{1\leq i\leq n})$ with imperfect information and winning objectives $(\varphi_i)_{1\leq i\leq n}$,

- a sequence of observations $\eta \in (O_i \times (\Sigma \cup \{\tau\}))^\omega$ is *doomsday compatible* (for Player i) if $\gamma_i(\eta) \subseteq \varphi_i \cup \bigcap_{j=1}^{j=n} \overline{\varphi_j}$, i.e. all plays that are compatible with η are either winning for Player i, or not winning for any other player,
- a prefix $\kappa \in (O_i \times (\Sigma \cup \{\tau\}))^* \cdot O_i$ of a sequence of observations is *good for retaliation* (for Player i) if there exists an observation-based strategy λ_i^R such that for all prefixes $\pi \in \gamma_i(\kappa)$ compatible with κ, $\mathsf{outcome}(\pi, \lambda_i^R) \subseteq \varphi_i \cup \bigcap_{j=1}^{j=n} \overline{\varphi_j}$.

The next lemma shows that the notions of sequences of observations that are doomsday compatible and good for retaliation prefixes are important for studying the existence of doomsday equilibria for imperfect information games.

Lemma 20. *Let G be an n-player game arena with imperfect information and winning objectives φ_i, $1 \leq i \leq n$. There exists a doomsday equilibrium in G if and only if there exists a play ρ in G such that:*

(F_1) *$\rho \in \bigcap_{i=1}^{i=n} \varphi_i$, i.e. ρ is winning for all the players,*

(F_2) *for all Player i, $1 \leq i \leq n$, for all prefixes κ of $\mathsf{Obs}_i(\rho)$, κ is good for retaliation for Player i,*

(F_3) *for all Player i, $1 \leq i \leq n$, $\mathsf{Obs}_i(\rho)$ is doomsday compatible for Player i.*

Proof. First, assume that conditions $(F_1), (F_2)$ and (F_3) hold and show that there exists a DE in G. We construct a DE $(\lambda_1, \ldots, \lambda_n)$ as follows. For each player i, the strategy λ_i plays according to the (observation of the) path ρ in $\mathcal{G}$, as long as the previous observations follow ρ. If an observation is unexpected for Player i (i.e., differs from the sequence in ρ), then λ_i switches to an observation-based retaliating strategy λ_i^R (we will show that such a strategy exists as a consequence of (F_2)). This is a well-defined profile and a DE because: (1) all strategies are observation-based, and the outcome of the profile is the path ρ that satisfies all objectives; (2) if no deviation from the observation of ρ is detected by Player i, then by condition (F_3) we know that if the outcome does not satisfy φ_i, then it does not satisfies φ_j, for all $1 \leq j \leq n$, (3) if a deviation from the observation of ρ is detected by Player i, then the sequence of observations of Player i so far can be decomposed as $\kappa = \kappa_1(o_1, \sigma_1) \ldots (o_m, \sigma_m)$ where (o_1, σ_1) is the first deviation of the observation of ρ, and (o_m, σ_m) is the first time it is Player i's turn to play after this deviation (so possibly $m = 1$). By condition (F_2), we know that κ_1 is good for retaliation. Clearly, $\kappa_1(o_1, \sigma_1) \ldots (o_\ell, \sigma_\ell)$ is retaliation compatible as well for all $\ell \in \{1, \ldots, m\}$ since retaliation goodness is preserved by player j's actions for all j. Therefore κ is good for retaliation and by definition of retaliation goodness there exists an observation-based retaliation strategy λ_i^R for Player i which ensures that regardless of the strategies of the opponents in coalition, if the outcome does not satisfy φ_i, then for all $j \in \{1, \ldots, n\}$, it does not satisfy φ_j either.

Second, assume that there exists a DE $(\lambda_1, \ldots, \lambda_n)$ in G, and show that $(F_1), (F_2)$ and (F_3) hold. Let ρ be the outcome of the profile $(\lambda_1, \ldots, \lambda_n)$. Then ρ satisfies (F_1) by definition of DE. Let us show that it also satisfies (F_3). By contradiction, if $\mathsf{obs}_i(\rho)$ is not doomsday compatible for Player i, then by definition, there is a path ρ' in $\mathsf{Plays}(G)$ that is compatible with the observations and actions of player i in ρ (i.e., $\mathsf{obs}_i(\rho) = \mathsf{obs}_i(\rho')$), but ρ' does not satisfy φ_i, while it satisfies φ_j for some $j \neq i$. Then, given the strategy λ_i from the profile, the other players in coalition can choose actions to construct the path ρ' (since ρ and ρ' are observationally equivalent for player i, the observation-based strategy λ_i is going to play the same actions as in ρ). This would show that the profile is not a DE, establishing a contradiction. Hence $\mathsf{obs}_i(\rho)$ is doomsday compatible for Player i for all $i = 1, \ldots, n$ and (F_3) holds. Let us show that ρ also satisfies (F_2). Assume that this not true. Assume that κ is a prefix of $\mathsf{obs}_i(\rho)$ such that κ is not good for retaliation for Player i for some i. By definition it means that the other players can make a coalition and enforce an outcome ρ', from some prefix of play compatible with κ, that is winning for one of players of the coalition, say Player j, $j \neq i$, and losing for Player i. This contradicts the fact that λ_i belongs to a DE. $\square$

We present automata constructions to recognise sequences of observations that are doomsday compatible and prefixes that are good for retaliation.

Lemma 21. *Given an n-player game G with imperfect information and a set of reachability, safety, parity, or LTL objectives $(\varphi_i)_{1 \leq i \leq n}$, we can construct for each Player i, a deterministic Streett automaton D_i whose language is exactly the set of sequences of observations $\eta \in (O_i \times (\Sigma \cup \{\tau\}))^\omega$ that are* doomsday compatible *for Player i, i.e.*

$$L(D_i) = \{\eta \in (O_i \times (\Sigma \cup \{\tau\}))^\omega \mid \forall \rho \in \gamma_i(\eta) \cdot \rho \in \varphi_i \cup \bigcap_{j \neq i} \overline{\varphi_j}\}.$$

For each D_i, the size of its set of states is bounded by $\mathbf{O}(2^{nk \log k})$ and the number of Streett pairs is bounded by $\mathbf{O}(nk^2)$ where k is the number of states in G for reachability, safety and parity objectives. For LTL *objectives, each D_i is at most doubly exponential in the size of the LTL objectives, and exponential in the size of the game arena and the number of players.*

Proof. Let $G = (S, (S_i)_{1\leq i\leq n}, s_{\mathsf{init}}, \Sigma, \Delta, (O_i)_{1\leq i\leq n})$, and let us show the constructions for Player i, $1 \leq i \leq n$. We treat the three types of winning conditions as follows.

We start with safety objectives. Assume that the safety objectives are defined implicitly by the following tuple of sets of safe states: $(T_1, T_2, \ldots, T_n)$, i.e. $\varphi_i = \mathsf{safe}(T_i)$. First, we construct the automaton

$$A = (Q^A, q^A_{\mathsf{init}}, (O_i \times (\Sigma \cup \{\tau\}), \delta^A)$$

over the alphabet $O_i \times (\Sigma \cup \{\tau\})$ as follows:

- $Q^A = S$, i.e. the states of A are the states of the game structure G,
- $q^A_{\mathsf{init}} = s_{\mathsf{init}}$,
- $(q, (o, \sigma), q') \in \delta^A$ if $q \in o$ and there exists $\sigma' \in \Sigma$ such that $\Delta(q, \sigma') = q'$ and such that $\sigma = \tau$ if $q \notin S_i$, and $\sigma = \sigma'$ if $q \in S_i$.

The acceptance condition of A is *universal* and expressed with LTL syntax:

A word w is accepted by A iff *all* runs ρ of A on w satisfy $\rho \models \Box T_i \vee \bigwedge_{j\neq i} \Diamond \overline{T_j}$.

Clearly, the language defined by A is exactly the set of sequences of observations $\eta \in (O_i \times (\Sigma \cup \{\tau\}))^\omega$ that are *doomsday compatible* for Player i, this is because the automaton A checks (using universal nondeterminism) that all plays that are compatible with a sequence of observations are doomsday compatible.

Let us show that we can construct a deterministic Streett automaton D_i that accepts the language of A and whose size is such that: (i) its number of states is at most $\mathbf{O}(2^{(nk \log k)})$ and (ii) its number of Streett pairs is at most $\mathbf{O}(nk)$. We obtain D with the following series of constructions:

- First, note that we can equivalently see A as the intersection of the languages of $n-1$ universal automata A_j with the acceptance condition $\Box T_i \vee \Diamond \overline{T_j}$, $j \neq i$, $1 \leq j \leq n$.
- Each A_j can be modified so that a violation of T_i is made permanent and a visit to $\overline{T_j}$ is recorded. For this, we use a state space which is equal to $Q^A \times \{0,1\} \times \{0,1\}$, the first bit records a visit to $\overline{T_i}$ and the second a visit to $\overline{T_j}$. We denote this automaton by A'_j, and its acceptance condition is now $\Box\Diamond(Q^A \times \{0,1\} \times \{0\}) \to \Box\Diamond(Q^A \times \{0\} \times \{0,1\})$. Clearly, this is a universal Streett automaton with a single Streett pair.
- A'_j, which is a universal Streett automaton, can be complemented (by duality) by interpreting it as a nondeterministic Rabin automaton (with one Rabin pair). This nondeterministic Rabin automaton can be made deterministic using a Safra like procedure, and according to [8] we obtain a deterministic Rabin automaton with $\mathbf{O}(2^{k \log k})$ states and $\mathbf{O}(k)$ Rabin pairs. Let us call this automaton A''_j.
- Now, A''_j can be complemented by considering its Rabin pairs as Streett pairs (by dualization of the acceptance condition): we obtain a deterministic Streett automaton with $\mathbf{O}(k)$ Streett pairs for each A_j.
- Now, we need to take the intersection of the $n-1$ deterministic automata A''_j (interpreted as Streett automata). Using a classical synchronized product we obtain a single deterministic Streett automaton D_i of size with $\mathbf{O}(2^{nk \log k})$ states and $\mathbf{O}(nk)$ Streett pairs. This finishes our proof for safety objectives.

Let us now consider reachability objectives. Therefore we now assume the states in $T_1, \ldots, T_n$ to be target states for each player respectively, i.e. $\varphi_i = \mathsf{reach}(T_i)$. The construction is in the same spirit as the construction for safety. Let $A = (Q^A, q^A_{\mathsf{init}}, O_i \times (\Sigma \cup \{\tau\}), \delta^A)$ be the automaton over $(O_i \times (\Sigma \cup \{\tau\})$ constructed from G as for safety, with the following (universal) acceptance condition;

A word w is accepted by A iff all runs ρ of A on w satisfy $\rho \models (\bigvee_{j\neq i} \Diamond T_j) \to \Diamond T_i$.

Clearly, the language defined by A is exactly the set of sequences of observations $\eta \in ((\Sigma \cup \{\tau\}) \times O_i)^\omega$ that are *doomsday compatible* for Player i (w.r.t. the reachability objectives). Let us show that we can construct a deterministic Streett automaton D_i that accepts the language of A and whose size is such that: (i) its number of states is at most $\mathbf{O}(2^{(nk \log k)})$ and (ii) its number of Streett pairs is at most $\mathbf{O}(nk)$. We obtain D_i with the following series of constructions:

- First, the acceptance condition can be rewritten as $\bigwedge_{j\neq i}(\Diamond T_j \to \Diamond T_i)$. Then clearly if A_j is a copy of A with acceptance condition $\Diamond T_j \to \Diamond T_i$ then $L(A) = \bigcap_{j\neq i} L(A_j)$.

- For each A_j, we construct a universal Streett automaton with one Streett pair by memorizing the visits to T_i and T_j and considering the acceptance condition $\Box\Diamond T_j \to \Box\Diamond T_i$. So, we get a universal automaton with a single Streett pair.

- Then we follow exactly the last three steps (3 to 5) of the construction for safety.

Let us now consider parity objectives. The construction is similar to the other cases. Specifically, we can take as acceptance condition for A the universal condition $\bigwedge_{j\neq i}(\mathsf{parity}_i \vee \overline{\mathsf{parity}_j})$, and treat each condition $\mathsf{parity}_i \vee \overline{\mathsf{parity}_j}$ separately. We dualize the acceptance condition of A, into the nondeterministic condition $\overline{\mathsf{parity}_i} \wedge \mathsf{parity}_j$. This acceptance condition can be equivalently expressed as a Streett condition with at most $\mathbf{O}(k)$ Streett pairs. This automaton accepts exactly the set of observation sequences that are not doomsday compatible for Player i against Player j. Now, using optimal procedure for determinization, we can obtain a deterministic Rabin automaton, with $\mathsf{O}(k^2)$ pairs that accepts the same language [31]. Now, by interpreting the pairs of the acceptance condition as Streett pairs instead of Rabin pairs, we obtain a deterministic Streett automaton A_j that accepts the set of observations sequences that are doomsday compatible for Player i against Player j. Now, it suffices to take the product of the $n-1$ deterministic Streett automata A_j to obtain the desired automaton A, its size is at most $\mathbf{O}(2^{nk \log k})$ with at most $\mathbf{O}(nk^2)$ Streett pairs.

For LTL objectives, we can proceed as follows. As for the other types of objectives, we construct for each player i from the arena G, a universal co-Büchi automaton on the alphabet $(O_i \times (\Sigma \cup \{\tau\}))$ with the following (universal) acceptance condition;

A word w is accepted by A iff all runs ρ of A on w satisfy $\rho \models \varphi_i \vee \wedge_{j\neq i, 1\leq j\leq n} \neg\varphi_j$.

Such an automaton has a state space exponential in the size of the LTL formulas and polynomial in the size of the arena. Then, we can turn this universal co-Büchi into a deterministic Streett automaton in exponential time, the state space of the resulting automaton is doubly exponential in the size of the LTL formulas, exponential in the size of the arena, and its number of Streett pairs is exponential in the size of the LTL formulas. □

Lemma 22. *Given an n-player game arena G with imperfect information and a set of reachability, safety or parity objectives $(\varphi_i)_{1\leq i\leq n}$, for each Player i, we can construct a finite-state automaton C_i that accepts exactly the prefixes of observation sequences that are good for retaliation for Player i.*

Proof. Let us show how to construct this finite-state automaton for any Player i, $1 \leq i \leq n$. Our construction follows these steps:

- First, we construct from G, according to lemma 21, a deterministic Streett automaton $D_i = (Q^{D_i}, q^{D_i}_{\mathsf{init}}, (O_i \times (\Sigma \cup \{\tau\}), \delta^{D_i}, \mathsf{St}^{D_i})$ that accepts exactly the set of sequences of observations $\eta \in (O_i \times (\Sigma \cup \{\tau\}))^\omega$ that are *doomsday compatible* for Player i.

- Second, we consider a turn-based game played on D_i by two players, A and B, that move a token from states to states along edges of D_i as follows:
  1. initially, the token is in some state q

2. then in each round: B chooses an observation $o \in O_i$ such that there exists $(q,(o,\sigma),q') \in \delta^{D_i}$. Then A chooses a transition $(q,(o,\sigma),q') \in \delta^{D_i}$ (which is completely determined by σ as D_i is deterministic), and the token is moved to q' where a new round starts.

The objective of A is to enforce from state q an infinite sequence of states, so a run of D_i that starts in q, and which satisfies St^{D_i} the Streett condition of D_i. For each q, this can be decided in time polynomial in the number of states in D_i and exponential in the number of Streett pairs in St^{D_i}, see [32] for an algorithm with the best known complexity. We denote by $\mathsf{Win} \subseteq Q^{D_i}$ the set of states q from which A can win the game above.

- Note that if $(o_1,\sigma_1)\dots(o_m,\sigma_m)$ is the trace of a path from q_{init} in D_i to a state $q \in \mathsf{Win}$, then clearly $(o_1,\sigma_1)\dots(o_{m-1},\sigma_{m-1})o_n$ is good for retaliation. Indeed, the winning strategy of A in q is an observation based retaliating strategy λ_i^R for Player i in G. On the other hand, if a prefix of observations reaches $q \notin \mathsf{Win}$ then by determinacy of Streett games, we know that B has a winning strategy in q and this winning strategy is a strategy for the coalition (against Player i) in G to enforce a play in which Player i does not win and at least one of the other players wins. So, from D_i and Win, we can construct a finite state automaton C_i which is obtained as a copy of D_i with the following acceptance condition: a prefix $\kappa = (o_0,\sigma_0),(o_1,\sigma_1),\dots,(o_{k-1},\sigma_{k-1}),o_k$ is accepted by C_i if there exists a path $q_0q_1\dots q_k$ in C_i such that q_0 is the initial state of C_i and there exists a transition labeled (o_k,σ) from q_k to a state of Win.

For reachability, safety and parity objectives, the overall complexity is exponential in the size of the game arena. This is because the state space of the Streett game is exponential in the size of the arena and the number of Streett pairs is polynomial in the size of the game arena.

For the LTL objectives, the overall complexity is doubly exponential in the size of the LTL formulas and exponential in the size of the game arena: this is because the state space of the game is doubly exponential in the size of the LTL formulas, exponential in the size of the game arena, and the number of Streett pairs is exponential in the size of the LTL formulas. □

Theorem 23. *The problem of deciding the existence of a doomsday equilibrium in an n-player game arena with imperfect information and n objectives is* ExpTime-C *for objectives that are either all reachability, all safety, all Büchi, all co-Büchi or all parity objectives, and is* 2ExpTime-C *for LTL objectives. Hardness already holds for 2-player game arenas.*

Proof. By Lemma 20, we know that we can decide the existence of a doomsday equilibrium by checking the existence of a play ρ in G that respects the conditions (F_1), (F_2), and (F_3). As we have shown in Lemma 22, for all $i \in \{1,\dots,n\}$, that the set of good for retaliation prefixes for Player i is definable by a finite-state automaton C_i, and the set of observation sequences that are doomsday compatible for Player i is definable by a Streett automaton D_i as we have shown in Lemma 21.

From the automata $(D_i)_{1\le i\le n}$ and $(C_i)_{1\le i\le n}$, we construct using a synchronized product a finite transition system T and check for the existence of a path in T that satisfy the winning objectives for each player in G, the Streett acceptance conditions of the $(D_i)_{1\le i\le n}$, and whose all prefixes are accepted by the automata $(C_i)_{1\le i\le n}$. The size of T is exponential in G for safety, reachability and parity objectives, and doubly exponential in the size of the objectives for LTL, and the acceptance condition is a conjunction of Streett and safety objectives. The existence of such a path can be established in polynomial time in the size of T, so in exponential time in the size of G for safety, reachability and parity objectives, and doubly exponential in the size of the objectives for LTL. The ExpTime-hardness for safety, reachability and parity objectives is a consequence of the ExpTime-hardness of two-player games of imperfect information for all the considered objectives [6, 10]. The 2ExpTime-hardness for LTL is a consequence of the 2ExpTime-hardness of the realizability problem for LTL [33]. □

5. Conclusion

We defined the notion of doomsday threatening equilibria both for perfect and imperfect information n player games with omega-regular objectives. This notion generalizes to n player games the winning secure equilibria [11]. Applications in the analysis of security protocols are envisioned and will be pursued as future works.

We have settled the exact complexity in games of perfect information for almost all omega-regular objectives with complexities ranging from PTime to 2ExpTime-C, the only small gap that remains is for parity objectives where we have a PSpace algorithm and both NP and coNP-hardness. For LTL, we also provide a Safraless solution [26] suitable to efficient implementation.

Surprisingly, the existence of doomsday threatening equilibria in n player games with imperfect information is decidable and more precisely ExpTime-C for safety, reachability and parity objectives, and 2ExpTime-C for LTL objectives.

References

[1] C. Schulte Althoff, W. Thomas, and N. Wallmeier. Observations on determinization of Büchi automata. *Theor. Comput. Sci.*, 363(2):224–233, 2006.

[2] R. Alur, T. A. Henzinger, and O. Kupferman. Alternating-time temporal logic. *Journal of the ACM*, 49:672–713, 2002.

[3] R. Alur and S. La Torre. Deterministic generators and games for LTL fragments. *TOCL*, 5, 2004.

[4] R. Alur, S. La Torre, and P. Madhusudan. Playing games with boxes and diamonds. In *CONCUR*, volume 2761 of *LNCS*, pages 127–141. Springer, 2003.

[5] Christel Baier and Joost-Pieter Katoen. *Principles of Model Checking (Representation and Mind Series)*. The MIT Press, 2008.

[6] D. Berwanger and L. Doyen. On the power of imperfect information. In *FSTTCS*, pages 73–82, 2008.

[7] J. R. Büchi and L. H. Landweber. Definability in the monadic second-order theory of successor. *J. Symb. Log.*, 34(2):166–170, 1969.

[8] Y. Cai, T. Zhang, and H. Luo. An improved lower bound for the complementation of Rabin automata. In *LICS*, pages 167–176. IEEE Computer Society, 2009.

[9] R. Chadha, S. Kremer, and A. Scedrov. Formal analysis of multiparty contract signing. *J. Autom. Reasoning*, 36(1-2):39–83, 2006.

[10] K. Chatterjee, L. Doyen, T.A. Henzinger, and J.-F. Raskin. Algorithms for omega-regular games with imperfect information. *LMCS*, 3(3), 2007.

[11] K. Chatterjee, T. A. Henzinger, and M. Jurdzinski. Games with secure equilibria. *Theor. Comput. Sci.*, 365(1-2):67–82, 2006.

[12] K. Chatterjee, T. A. Henzinger, and N. Piterman. Strategy logic. *Inf. Comput.*, 208(6):677–693, 2010.

[13] K. Chatterjee, T.A. Henzinger, and N. Piterman. Generalized parity games. In *FoSSaCS*, volume 4423 of *LNCS*, pages 153–167. Springer, 2007.

[14] A. Da Costa Lopes, F. Laroussinie, and N. Markey. ATL with strategy contexts: Expressiveness and model checking. In *FSTTCS*, volume 8 of *LIPIcs*, pages 120–132, 2010.

[15] E. A. Emerson and C. Jutla. Tree automata, mu-calculus and determinacy. In *FOCS*, pages 368–377. IEEE Comp. Soc., 1991.

[16] E.A. Emerson and C.-L. Lei. Modalities for model checking: Branching time strikes back. In *POPL*, pages 84–96, 1985.

[17] N. Fijalkow and F. Horn. The surprizing complexity of reachability games. *CoRR*, abs/1010.2420, 2010.

[18] Emmanuel Filiot, Naiyong Jin, and Jean-François Raskin. An antichain algorithm for LTL realizability. In *CAV 2009*, volume 5643 of *Lecture Notes in Computer Science*, pages 263–277. Springer, 2009.

[19] Emmanuel Filiot, Naiyong Jin, and Jean-François Raskin. Antichains and compositional algorithms for LTL synthesis. *Formal Methods in System Design*, 39(3):261–296, 2011.

[20] Bernd Finkbeiner and Sven Schewe. Bounded synthesis. *STTT*, 15(5-6):519–539, 2013.

[21] D. Fisman, O. Kupferman, and Y. Lustig. Rational synthesis. In *Proc. of TACAS*, LNCS 6015, pages 190–204. Springer, 2010.

[22] E. Grädel, W. Thomas, and T. Wilke, editors. *Automata, Logics, and Infinite Games: A Guide to Current Research*, LNCS 2500. Springer, 2002.

[23] N. Immerman. Number of quantifiers is better than number of tape cells. *Journal of Computer and System Sciences*, 22:384–406, 1981.

[24] W. Jamroga, S. Mauw, and M. Melissen. Fairness in non-repudiation protocols. In *Proc. of STM*, LNCS 7170, pages 122–139. Springer, 2012.

[25] S. Kremer and J.-F. Raskin. A game-based verification of non-repudiation and fair exchange protocols. *Journal of Computer Security*, 11(3):399–430, 2003.

[26] O. Kupferman and M. Y. Vardi. Safraless decision procedures. In *FOCS*, 2005.

[27] D. Martin. Borel determinacy. In *Annals of Mathematics*, volume 102, pages 363–371, 1975.

[28] F. Mogavero, A. Murano, G. Perelli, and M.Y. Vardi. What makes ATL* decidable? a decidable fragment of strategy logic. In *CONCUR*, pages 193–208, 2012.

[29] F. Mogavero, A. Murano, and M. Y. Vardi. Reasoning about strategies. In *Proc. of FSTTCS*, volume 8 of *LIPIcs*, pages 133–144. Schloss Dagstuhl - LZfI, 2010.

[30] J. F. Nash. Equilibrium points in n-person games. *PNAS*, 36:48–49, 1950.

[31] N. Piterman. From nondeterministic Büchi and Streett automata to deterministic parity automata. *Logical Methods in Computer Science*, 3(3), 2007.

[32] N. Piterman and A. Pnueli. Faster solutions of Rabin and Streett games. In *LICS*, pages 275–284, 2006.

[33] A. Pnueli and R. Rosner. On the synthesis of a reactive module. In *POPL*, pages 179–190. ACM Press, 1989.

[34] M. O. Rabin. Decidability of second-order theories and automata on infinite trees. *Trans. Amer. Math. soc.*, 141:1–35, 1969.

[35] P. J. Ramadge and W. M. Wonham. Supervisory control of a class of discrete event processes. *SIAM Journal on Control and Optimization*, 25(1):206–230, 1987.

[36] Shmuel Safra. On the complexity of omega-automata. In *29th Annual Symposium on Foundations of Computer Science, White Plains, New York, USA, 24-26 October 1988*, pages 319–327. IEEE Computer Society, 1988.

[37] L. S. Shapley. Stochastic games. *PNAS*, 39:1095–1100, 1953.

[38] M. Ummels and D. Wojtczak. The complexity of Nash equilibria in stochastic multiplayer games. *Logical Methods in Computer Science*, 7(3), 2011.

[39] F. Wang, C. Huang, and F. Yu. A temporal logic for the interaction of strategies. In *Proc. of CONCUR*, LNCS 6901, pages 466–481. Springer, 2011.